
MIND: MARKET INTERPRETATION DSL FOR UNIFIED MARKET DESIGN AND SIMULATION

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 Market mechanisms such as auctions and matchings coordinate supply and de-
012 mand at scale, yet their implementations remain locked in rigid procedural code
013 that hinders iteration and auditing. We introduce the Market Interpretation DSL
014 (MIND), a typed language and toolchain for declarative market specification to
015 achieve unified market design and simulation. MIND comprises (i) a core gram-
016 mar with a phased Intermediate Representation (IR) and economic safety checks,
017 (ii) a natural language assistant that translates descriptions into DSL with auto-
018 mated diagnostics and safe rewrites, and (iii) rule-based simulation and convex op-
019 timization backends. Using synthetic specifications generated across 87 domains
020 with held-out validation, our fine-tuned Llama-3-8B assistant achieves 96.33%
021 semantic correctness, measured as IR equivalence to gold programs, surpassing
022 few-shot GPT-4o at 81.11%. Across second-price auctions, multi-stage auctions,
023 and matching markets, MIND reduces specification complexity by approximately
024 79% in lines of code compared to Python implementations. In a within-subjects
025 study with 17 participants, mechanism modifications were completed 4 to 10 times
026 faster using MIND. Code, dataset, and models will be released upon acceptance.
027

1 INTRODUCTION

028 Market mechanisms such as auctions and matching markets form the backbone of modern eco-
029 nomics, digital platforms, and decentralized systems. They coordinate supply and demand, reduce
030 transaction costs, and enable efficient allocation of scarce resources (Milgrom, 2004; Roth, 2018;
031 Milgrom, 2021). Despite this centrality, practical modeling and implementation remain cumber-
032 some. Most platforms and simulators still hard-code allocation rules, matching logic, and pricing
033 routines into procedural code, creating a lossy translation from policy to code (Calheiros et al., 2011;
034 Byrd et al., 2019). This limits experimentation and complicates verification of market properties.
035

036 Beyond performance, platform operators must ensure transparent rules and reproducible outcomes
037 for regulatory compliance, requiring a chain from human-readable policies to executable logic with
038 audit trails. Decentralized trading systems raise the bar further: mechanism logic executes on-chain
039 and requires formal checks for correctness and economic safety (d’Eon et al., 2024; Bouaicha et al.,
040 2025). The core limitation is the absence of a unified interface that bridges conceptual specifications
041 to deterministic implementations with support for governance, testing, and audit.

042 Recent LLM-based approaches to mechanism automation yield non-deterministic outputs and brit-
043 tle patches, making debugging difficult where fairness and correctness are paramount. They also
044 struggle to bridge underspecified natural language and verbose implementations, frequently omit-
045 ting crucial details such as reserve prices, tie-breaking rules, or budget constraints. Moreover, the
046 primary users are economists and policy analysts who possess domain expertise but typically lack
047 programming skills. While GUI-based tools exist, they cannot express conditional constraints or
048 multi-stage interactions, reducing mechanisms to rigid templates.

049 We advocate a domain-specific language paired with natural language translation that separates au-
050 thoring, validation, and execution. A compact, typed DSL makes specifications legible, enables
051 static checks for economic consistency, and supports deterministic compilation to multiple back-
052 ends for cross-validation. This creates a governance surface where specifications carry provenance,
053 version identifiers, and audit traces, while validators enforce safety gates before deployment. The
system serves both audiences: domain experts use natural language to generate initial specifications,

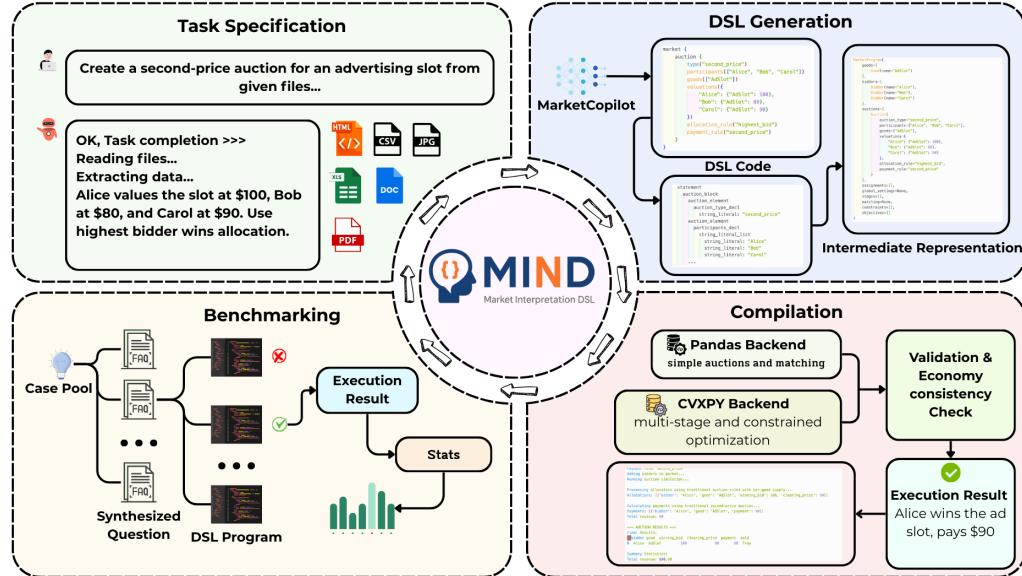


Figure 1: The whole workflow of our system. First, the Completeness Agent helps users with complete task descriptions. Then, the [AI Copilot \(MarketCopilot\)](#) translates the refined description into MIND programs. Finally, the compiler executes the programs with appropriate backends to produce the final results.

while power users can directly edit DSL programs. In MIND, semantic correctness is evaluated as Intermediate Representation (IR) equivalence to reference programs, and specification complexity is measured in AST nodes.

In this paper, we present MIND (Market Interpretation DSL), a comprehensive and extensible language and toolchain for specifying, validating, and executing market mechanisms. MIND defines a grammar of objects, actions, and types, compiled into a phased IR which is automatically validated and compiled into executable simulations, thereby decoupling specification from execution. The system includes a natural language assistant that translates descriptions into DSL programs, provides diagnostics, and applies safe rewrites (Zhang et al., 2023). Two backends support rule-based simulation and convex optimization. Each specification is a versioned artifact with machine-checkable metadata.

We evaluate MIND along three axes. First, a workflow study shows a 79% reduction in specification complexity versus Python. Second, a within-subjects user study ($N = 17$) confirms that practitioners modify mechanisms $4\text{--}10\times$ faster in MIND than in a Python baseline. Third, we demonstrate generalization by encoding 8 mechanisms from recent ACM EC/SIGecom literature without changing the core grammar. Additionally, our fine-tuned Llama-3-8B Copilot achieves 96.33% semantic correctness on NL \rightarrow DSL translation.

Our contributions are threefold:

- We introduce MIND, a domain-specific language with formal grammar and a phased IR that bridges natural language to executable simulations while creating auditable specifications.
- We develop an execution framework with two backends that support rule-based simulation and convex optimization from unified specifications, enabling deterministic cross-validation for compliance.
- We build a natural language to DSL translation system, achieving 96.33% semantic correctness across 87 domains, serving both non-programmers and power users.

2 RELATED WORK

Market mechanism DSLs. Prior mechanism modeling in economics has largely relied on general-purpose programming or specialized simulators, making specification and validation cumbersome. Recent work explores domain-specific languages to capture auction rules, negotiation

108 games, or fair division protocols in symbolic form (Hoseindoost et al., 2024; De Jonge & Zhang, 2021; Bertram et al., 2023). CoorERE (Hoseindoost et al., 2024) provides an executable DSL for 109 auction-based coordination in crisis response, reducing development effort by nearly half, but it 110 addresses single-item auctions without cross-mechanism support. GDL has been repurposed as a 111 unifying description language for negotiation domains (De Jonge & Zhang, 2021), enabling generic 112 solvers, but it lacks intermediate representations with economic validation. Slice (Bertram et al., 113 2023) defines a DSL for fair division protocols with automated envy-freeness verification, yet re- 114 mains limited to division problems without auction or matching support. These DSLs improve 115 mechanism specification but are scoped to individual subdomains and do not provide staged valida- 116 tion, two execution backends, or governance artifacts that MIND includes. 117

118 **LLMs in mechanism design.** The rise of large language models has motivated new approaches 119 to automating specification and simulation. Recent studies use LLMs to generate valuations, bid- 120 ding policies, and to propose new auction formats (Duetting et al., 2024; Sun et al., 2024; Dubey 121 et al., 2024b; Shah et al., 2025). LaMP-Val (Sun et al., 2024) uses GPT-4 to infer personalized 122 valuations from text and fine-tunes smaller models as strategic agents. Dubey et al. (Dubey et al., 123 2024b) and Duetting et al. (Duetting et al., 2024) examine auctions where advertisers bid for in- 124 fluence over LLM outputs, proposing incentive-compatible rules for token-level allocation. Shah et 125 al. (Shah et al., 2025) show GPT-4 agents can reproduce human-like bidding behaviors, suggesting 126 LLMs can serve as synthetic participants. These approaches demonstrate potential for synthesis and 127 simulation but operate without typed specifications, deterministic compilation, or audit trails. They 128 generate code directly without an intermediate representation, making systematic verification and 129 governance difficult. They often lack empirical validation of generated mechanisms against ground 130 truth specifications. 131

131 **Unified frameworks and positioning.** Prior DSLs achieve domain-specific expressiveness and 132 LLM approaches enable automation, yet the literature remains fragmented: CoorERE focuses on 133 crisis response, Slice on fair division, and LLM methods typically lack formal specifications. Tech- 134 nical barriers to unification include incompatible type systems across auction and matching domains, 135 the absence of staged validation for economic properties, and limited support for governance re- 136 quirements such as provenance tracking and policy diffs. MIND addresses these gaps through a 137 unified grammar spanning auctions, matchings, and exchanges; an intermediate representation with 138 three-stage validation (parsing, typing, economic consistency); two execution backends for simula- 139 tion and optimization that scale to thousands of participants; natural language translation achieving 140 96.33% semantic correctness on 87 domains; and governance artifacts including versioning, 141 validator reports, and audit logs. This combination links formal specification, property verification, 142 and agent-based evaluation in a single reproducible workflow. Our evaluation shows it reduces 143 specification complexity by 79% while maintaining semantic accuracy comparable to hand-written 144 implementations. 145

145 **Verification and Guardrail Systems.** Beyond synthesis, ensuring safety requires rigorous ver- 146 ification. Formal verification engines like Imandra (Passmore et al., 2020) and certified auction 147 frameworks (Caminati et al., 2015) use theorem proving to guarantee properties like incentive com- 148 patibility. In industry, frameworks like AWS Bedrock Guardrails enforce output safety policies. 149 MIND complements these systems rather than replacing them: by producing a typed, distinct Inter- 150 mediate Representation (IR) rather than opaque Python scripts, MIND provides the necessary 151 structured input that these advanced verification engines require to perform mathematical proofs 152 and policy enforcement. 153

154 3 METHOD

155 As illustrated in Figure 1, our system provides an end-to-end pipeline for generating, validating, 156 and simulating MIND, starting from a user specification. The architecture is composed of several 157 key parts: (1) a symbolic DSL for formal representation (Shi et al., 2024; Borum & Seidl, 2022), 158 (2) an Intermediate Representation (IR) with a robust validation system, (3) a two-backend frame- 159 work for code generation, and (4) an AI-powered toolchain including a dataset generation pipeline, 160 a completeness agent, and a [fine-tuned AI Copilot](#) (referred to as [MarketCopilot](#)) for [NL](#)→[DSL](#) 161 translation. 162

162 3.1 MARKET INTERPRETATION DSL
163

164 The foundation of our system is Market Interpretation DSL (MIND), a formal language designed
165 for the specification of market rules. The language’s grammar is built on a clear separation of core
166 concepts: (i) **Objects** are entities that constitute a market, such as auction, participants, goods, and
167 matching; (ii) **Actions** are operations that define the market’s behavior, such as specifying the auc-
168 tion type, defining valuations, or setting constraints; (iii) **Types** are specific variants of objects and
169 actions, like `type("second_price")` or `type("first_price")`. Some DSL examples
170 are shown in Figure 2.

171 **Scope of Expressiveness** From the current grammar and IR, MIND is designed to support:
172 (1) single- and multi-stage sealed-bid auctions with standard allocation and payment rules; (2)
173 compatibility-constrained matching markets (e.g., stable matching (Gale & Shapley, 1962)) with
174 explicit compatibility graphs; and (3) mechanisms with convex objectives and linear constraints.
175 We explicitly detail the supported features versus out-of-scope capabilities in Appendix Table 4.
176 Note that we do not currently support fully general combinatorial bidding languages (e.g., arbitrary
177 XOR bids without constraints (Nisan, 2006)) or iterative open-cry auctions (e.g., ascending clock
178 auctions (Ausubel, 2004)); these are planned for future extensions.

179 3.2 INTERMEDIATE REPRESENTATION (IR) AND VALIDATION
180

181 A challenge in designing a language with multiple execution targets is preventing our language
182 parser from getting entangled with the specific details of every execution backend (Pandas, CVXPY,
183 etc.). This creates a brittle, unscalable system where adding a new backend or modifying the DSL
184 syntax would require cascading changes across the entire codebase.

185 To solve this problem, we introduce an Intermediate Representation (IR) (Lattner et al., 2021) as a
186 critical abstraction layer. The IR serves as the **single source of truth** for mechanism semantics and
187 governance, not just an intermediate parsing artifact. It is a typed abstract syntax tree (AST) over
188 market constructs (e.g., `AuctionNode`, `ConstraintNode`). The parser translates DSL source
189 into IR only; code generators read IR only. This separation ensures modularity and maintainability.
190 Crucially, IR nodes, validator report identifiers, and compile artifacts are logged with each execution,
191 so any result can be replayed with the exact spec and validator configuration.

192 To ensure that any market specified in the DSL is not just syntactically correct, but also seman-
193 tically and economically sound, the IR undergoes a rigorous validation process before generating
194 code. This process consists of three phases: parsing, typing, and economic consistency. Three main
195 validators run in order on the IR:

1. **CoreMarketValidator:** Performs fundamental checks, ensuring names are unique, references
197 are valid, valuations align with participants, and auction rules are recognized.
2. **StageAndMatchingValidator:** If the design uses stages or matching, this validator runs to per-
199 form checks on global settings and validate the structure of these advanced components.
3. **AdvancedOptimizationValidator:** If the design includes constraints or objectives, this validator
201 checks that their types are recognized and parameters are valid (e.g., non-negative budgets).

203 Each validator consumes an IR snapshot and emits a `ValidationReport` with typed findings
204 (`error`, `warning`, `autofixable`). The Autofixer applies only safe rewrites; if a required rule
205 cannot be inferred, it emits a blocking error rather than altering semantics. We persist the spec hash,
206 validator report identifier, and compile artifact path with the run logs to enable exact reconstruction
207 during audit. All experiments log these identifiers, allowing any reported result to be traced to its
208 exact specification and validator state.

209 3.3 TWO-BACKEND CODE GENERATION
210

211 While theoretically a monolithic engine could handle all designs, we treat backend selection as
212 an engineering trade-off between performance and solver generality. Simple single-shot auctions
213 and matching markets benefit from fast, vectorized table-driven simulation, whereas constrained or
214 combinatorial designs require solver-grade convex optimization. To handle both without exposing
215 backend complexity to users, we compile the same **backend-independent IR**—the **single source of**
216 **truth**—into different execution targets via `MarketCompiler`.

Figure 2: Two MIND specifications. Left: second-price auction where `type` specifies auction type, `participants` lists bidders, `goods` declares the item, `valuations` gives each bidder's valuation, `allocation_rule()` assigns to the highest bidder, and `payment_rule()` charges the second-highest bid. Right: simple matching market specification.

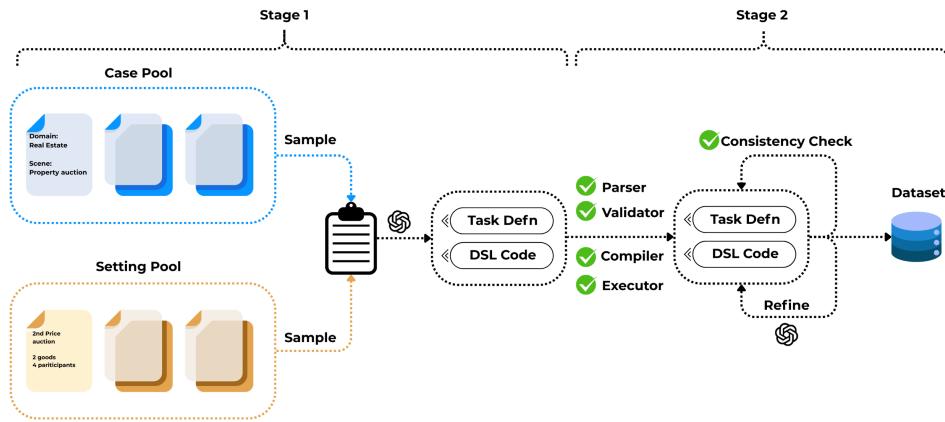


Figure 3: Market Mechanism Dataset pipeline with three phases: data generation, augmentation, and filtering.

Backend routing MarketCompiler selects a backend via a deterministic function of IR features: designs without explicit objectives or global constraints are routed to the **Pandas/NetworkX (Simulation)** backend; those declaring objectives or linear constraints are routed to the **CVXPY (Optimization)** backend. Crucially, all backends must implement identical IR semantics; observable outcomes (allocations, payments) must agree for any given specification, enabling cross-backend consistency checks for governance.

- **Pandas Backend:** A lightweight, simulation-based backend for standard auctions and matching markets, optimized for speed and simplicity.
- **CVXPY Backend:** An optimization-based backend using the CVXPY library (Diamond & Boyd, 2016), automatically selected for scenarios involving constraints (e.g., budget balance) or objectives (e.g., maximizing revenue).

Backends implement the same IR semantics; observable outcomes (allocations, payments, feasibility flags) must agree for identical IR inputs. This two-backend approach lets users scale from simple simulations to constrained optimizations without changing the DSL specification, while preserving consistent semantics across backends.

3.4 DATASET GENERATION

Supervised fine-tuning of a specialized Copilot requires a large-scale, high-quality dataset of (Natural Language Description, DSL) pairs. To the best of our knowledge, there is no such dataset for the task of translating natural language specifications into a formal DSL for market mechanisms. To

270 address this, we developed an automated pipeline (Ratner et al., 2017; Northcutt et al., 2021) to use
271 LLMs to generate synthetic data as illustrated in Figure 3. The process begins by programmatically
272 generating diverse prompts for a generator LLM (GPT-4o (Hurst et al., 2024)). We predefine over
273 **900 possible market use cases within 87 domains**. For each use case, we randomly sample settings
274 to generate prompts with the formal DSL grammar and in-context examples.

275 To ensure correctness, every DSL program undergoes a rigorous **4-stage validation pipeline**: (1)
276 **grammar parsing**, (2) **three-phase IR validation (structure, typing, economic consistency)**, (3) **com-**
277 **ilation to both backends**, and (4) **execution without runtime errors**.

278 After guaranteeing code correctness, we refine the corresponding natural language descriptions.
279 Each validated DSL program is passed to an LLM to generate a more detailed and complete de-
280 scription. As a quality control step, another verifier LLM performs a description-DSL consistency
281 check, confirming semantic alignment between the enhanced natural language description and the
282 DSL code. **Finally, beyond these automated checks, we conducted a manual inspection of the can-**
283 **didate pool to filter out any remaining low-quality or redundant samples**. Only pairs passing this
284 multi-stage verification are included in the dataset.

285 3.4.1 HUMAN AUDIT PROTOCOL

286 To ensure the quality of our automated pipeline, we perform a manual human audit. We drew a
287 simple random sample of 100 (description, DSL) pairs from the final, post-filter dataset after the 4-
288 stage validation and description-DSL consistency check, stratified by domain and mechanism type.
289 An expert evaluator (one of the authors) assessed each pair on: (i) syntactic correctness (DSL parses
290 under the grammar), (ii) semantic alignment (IR equivalence of the compiled DSL to the behavior
291 described), and (iii) functional executability (successful backend compilation and execution). The
292 evaluator was blinded to the specific pipeline metadata during assessment to minimize bias. A pair
293 passes only if all three criteria are satisfied. The audit confirms a very high accuracy rate under these
294 criteria.

295 3.5 COMPLETENESS AGENT

296 To ensure the generated DSL program fully aligns with users’ expectations, we need descriptions
297 with sufficient detail. Since users rarely provide complete descriptions initially, we developed a
298 Completeness Agent (Yao et al., 2023; Shinn et al., 2023) as a pre-processor to help users provide
299 enough information for the MarketCopilot.

300 The agent operates through a multi-node scheme defining key elements to capture. For each node, it
301 extracts information from users’ prompts and populates the scheme with required and optional fields.
302 If all required fields are fulfilled, the node completes and the agent proceeds. Otherwise, it requests
303 missing critical information. After passing all nodes, the agent outputs a *completion schema* and
304 a *complete task specification* that are passed to the MarketCopilot. The completion schema maps
305 directly to the MarketCopilot input fields (participants, goods, constraints, objectives), ensuring the
306 prompt is structurally complete before translation. This significantly increases the likelihood of
307 generating a valid and executable DSL program on the first attempt.

308 3.6 MARKETCOPILOT FINETUNING

309 The core of our natural language interface is the **AI Copilot (MarketCopilot)**, an AI assistant
310 fine-tuned for NL-to-DSL translation. It is distinct from the Completeness Agent (which only fills
311 missing fields) and the DSL framework itself. We used a Llama-3-8B-Instruct model (Dubey et al.,
312 2024a) as our base and applied LoRA (Hu et al., 2022) with rank $r = 32$ for efficient training. The
313 model was trained on our curated dataset using a standard supervised fine-tuning (SFT) objective to
314 maximize the conditional probability of generating the ground-truth DSL program given the natural
315 language description. The training loss is:

$$316 \mathcal{L}(\theta) = - \sum_{(X_i, Y_i) \in \mathcal{D}} \log P(Y_i | X_i; \theta) \quad (1)$$

317 Here, \mathcal{D} is the set of (description X_i , DSL sequence Y_i) pairs; Y_i is tokenized as a left-to-right
318 sequence for teacher forcing under the SFT objective in Eq. 1.

324
325

327 3.7 PLUGGABLE FRONTEND INTERFACES

328 We view natural-language and graphical user interfaces (GUIs) as complementary rather than competing approaches. The MIND architecture is designed with a **pluggable front-end**: the core DSL and IR remain invariant regardless of the input modality. To validate this, we developed a **prototype web GUI** where users input natural language descriptions, and the Completeness Agent runs interactively to ask follow-up questions and populate a structured schema before passing it to the MarketCopilot. This demonstrates that MIND supports diverse workflows—whether pure natural language, form-based, or hybrid—all targeting the same unified, verifiable DSL and IR.

333

334 4 EXPERIMENTS

335

336 To validate our system, we designed two primary experiments. First, we evaluate the Market Interpretation DSL itself by comparing its workflow and expressiveness against standard procedural 337 programming approaches for market design. Second, we quantitatively evaluate the performance of 338 our fine-tuned AI Copilot in translating natural language specifications into valid and semantically 339 correct DSL code.

340

341 4.1 DSL WORKFLOW AND USER STUDY EVALUATION

342

343 The primary motivation for creating a Domain-Specific Language is to accelerate development, and 344 reduce errors. We evaluate MIND’s impact on market design workflows at two levels: (i) a **code-level 345 comparison of hand-written implementations** (Table 1) to measure specification complexity; and (ii) a **within-subjects user study** with 17 participants comparing MIND against a Python baseline 346 to measure modification efficiency.

347

348 4.1.1 METHODOLOGY AND BENCHMARK TASKS

349

350 **Code-Level Benchmarks (Table 1).** To quantify specification complexity, we implemented three 351 representative market mechanisms: (1) A standard second-price auction; (2) A multi-stage auction 352 with reserves; and (3) A compatibility matching market. We implemented each using MIND, standard 353 Python, and **AnyLogic** (Borshchev & Filippov, 2004). We evaluated these approaches based 354 on three established software engineering criteria. **Specification Complexity** refers to the effort 355 required to define the mechanism, measured by the number of distinct modeling steps and Source 356 Lines of Code (LoC) (Molnar & Motogna, 2020). **Readability & Verifiability** describes how easily 357 the implementation can be audited against its theoretical design, a crucial aspect of model correctness 358 (Alawad et al., 2019). Finally, **Flexibility** measures the effort required to modify an existing 359 mechanism (e.g., changing a pricing rule), which is a key indicator of software maintainability 360 (Ardito et al., 2020).

361

362 **User Study Methodology.** We recruited 17 participants with varying levels of expertise. The 363 majority had >2 years of Python experience. The study followed a within-subjects design where each 364 participant performed modification tasks on the three benchmark mechanisms in both MIND and 365 Python environments. The tasks required modifying pricing rules, adding re-auction stages, and 366 adjusting compatibility constraints. We recorded Completion Time (self-reported minutes) and Correctness 367 (verified offline). Detailed protocol and participant demographics are provided in **Appendix I**.

368

369 4.1.2 RESULTS AND ANALYSIS

370

371 **Specification Complexity** (Table 1). The results of our code-level comparison demonstrate the 372 significant advantages of the DSL-based approach. As summarized in Table 1, MIND consistently 373 requires the least specification effort. Across all three tasks, MIND reduces specification complexity 374 by approximately 79% in lines of code compared to Python implementations. Its declarative syntax 375 allows designers to focus on economic rules rather than procedural control flow. In contrast, the 376 Python approach requires significant boilerplate code and embeds the core mechanism logic within 377 procedural control flow, making it difficult to verify and modify. While powerful, AnyLogic introduced a high degree of complexity and a steep learning curve, making it not that suits for the rapid 378 prototyping of mechanism rules, which is a primary goal of our system.

379

380 **User Study Results.** The results of our user study indicate a significant workflow advantage for the 381 DSL. Participants completed the mechanism modifications consistently faster using MIND com-

378
 379 Table 1: A comparison of implementation workflows for common market design tasks. Our declar-
 380 ative approach, MIND, consistently requires the least specification effort, offers the highest clarity
 381 for verification, and provides the greatest flexibility for experimentation.

Task	Approach	Specification Complexity	Readability & Verifiability	Flexibility (Effort to Modify)
Second-Price Auction	MIND	~10 lines of DSL	High: Declarative economic syntax.	Trivial: Change a single keyword.
	Python	~40-60 lines of code	Low: Core logic is embedded in code.	Moderate: Requires rewriting functions.
	AnyLogic	~15-20 graphical steps	Medium: Logic is distributed across agents.	High: Requires reconfiguration.
Multi-Stage Auction	MIND	~20-25 lines of DSL	High: Staging logic is explicit and easy to follow.	Trivial: Modify a self-contained stage block.
	Python	~100-120 lines of code	Low: State management is complex and error-prone.	High: Requires significant refactoring of the main control flow.
	AnyLogic	~25-35 graphical steps	Low: Managing agent state across stages is hard.	High: Requires a full redesign of the simulation flowchart.
Matching Market	MIND	~15 lines of DSL	High: The compatibility graph is a direct input.	Low: Change the data directly.
	Python (w/ NetworkX)	~60-80 lines of code	Medium: Requires graph library expertise to understand.	Moderate: Requires implementing a different matching algorithm.
	AnyLogic	~35-45 graphical steps	Medium: Requires defining custom agent interaction rules.	High: Requires creating new agent protocols.

pared to Python (with median speed-up between 4-10x). Detailed participant demographics, task breakdowns, and full statistical reports are provided in [Appendix I](#).

4.2 COPILOT GENERATION EVALUATION

This experiment evaluates the ability of our [AI Copilot \(MarketCopilot\)](#) to automatically generate high-quality Market Interpretation DSL code from natural language descriptions.

4.2.1 EXPERIMENTAL SETUP

We evaluated our fine-tuned MarketCopilot (Llama-3-8B + LoRA) against several baseline models on a held-out test set of 323 examples across 87 domains, ensuring no overlap with training data beyond a 0.85 cosine similarity threshold computed on TF-IDF representations. We additionally exclude near-duplicates by AST hash to prevent leakage. The training set consisted of 11,000 examples, with 10% used for validation during hyperparameter tuning.

To provide a robust comparison, baseline models were evaluated in a few-shot setting with the formal DSL grammar specification and 4 examples of (NL, DSL) pairs along with the task description. In contrast, our MarketCopilot operates zero-shot, taking only the natural language task description as input.

4.2.2 EVALUATION PIPELINE AND METRICS

We employed a rigorous multi-stage pipeline to assess correctness:

1. **Grammar Validation:** Each output is parsed using our Lark EBNF grammar. We measure Parse Success Rate as the percentage of syntactically valid programs.
2. **Semantic Validation:** Syntactically correct programs undergo three checks:
 - *Validator check:* Tests logical consistency using the three-phase validation (parsing, typing, economic consistency)
 - *Compiler check:* Verifies code generation to the two execution backends

432
 433 Table 2: Performance of the AI Copilot against baseline LLMs on the NL-to-DSL generation task
 434 (323 test cases). **Our fine-tuned model outperforms significantly larger proprietary models**,
 435 demonstrating that strict grammar compliance requires domain adaptation rather than just reasoning
 436 power.

437 Model	Parse Success (%)	Val + Comp Success (%)	IR Equivalence (%)
<i>Proprietary SOTA Models (Few-shot)</i>			
GPT-4o-mini	97.52	89.47	76.78
GPT-4o (Hurst et al., 2024)	98.45	95.98	81.11
Gemini 2.5 Pro (Reasoning)	99.69	99.07	77.71
GPT-5 (Reasoning)	96.90	92.57	72.12
<i>Open-Source Models</i>			
Llama-3-8B	80.76	73.71	60.89
Qwen3-Coder-30B(Yang et al., 2025)	95.51	94.55	81.73
Our Method (Llama-3 + LoRA)	100.00	100.00	96.33

442
 443 • *IR Semantic Equivalence*: We define this as graph isomorphism between the generated and
 444 reference IR ASTs, modulo variable renaming, commutative ordering, and alias normalization
 445

446 3. **Execution Validation:** Programs are executed on 323 test scenarios to verify they produce cor-
 447 rect market outcomes (allocations, payments, feasibility). Scenarios mirror the functional spec
 448 used in the workflow study.

449 Our primary metric is **End-to-End Correctness**: the percentage of generations that pass all vali-
 450 dation stages and are semantically equivalent to the reference solution. We define this strictly: a
 451 generation is correct if and only if it: (1) parses successfully, (2) passes all three validation stages,
 452 (3) compiles successfully to both backends, (4) achieves IR semantic equivalence to the ground
 453 truth, and (5) passes execution checks.

454
 455 4.2.3 RESULTS AND DISCUSSION

456 Table 2 presents the performance metrics. Our fine-tuned model achieves the highest scores across
 457 all metrics, with **96.33% end-to-end correctness**, significantly outperforming the best general-
 458 purpose baseline (GPT-4o at 81.11%).

459 Crucially, our experiments reveal that stronger reasoning models do not necessarily yield better DSL
 460 specifications. The reasoning-focused **GPT-5** model achieves only 72.12% correctness, performing
 461 worse than GPT-4o. Error analysis indicates that powerful reasoning models tend to “over-think” the
 462 task: they frequently hallucinate plausible but unsupported keywords or attempt to restructure the
 463 mechanism logic in ways that violate the DSL schema. This highlights a key trade-off: while large
 464 models excel at unstructured reasoning, domain-specific fine-tuning is essential for strict adherence
 465 to formal grammars.

466 It is also worth noting that MIND is **frontend-agnostic**. While our current implementation uses a
 467 fine-tuned 8B model for efficiency and controllability, the underlying IR, validator, and compiler
 468 pipeline can serve as a robust guardrail for any future foundation model.

469 4.3 GENERALIZATION TO ACADEMIC MECHANISMS

470 To assess MIND’s capability to generalize beyond standard textbook examples, we implemented
 471 8 mechanisms from recent ACM EC and SIGecom papers (2020–2024). These include complex
 472 designs such as dynamic auction throttling, randomized clock auctions, and stochastic ridesharing
 473 matching.

474 **Results.** All 8 mechanisms were successfully encoded in MIND, passed the 4-stage validation
 475 pipeline, and executed on the appropriate backend (simulation or optimization). Crucially, none
 476 required changes to the core MIND grammar; specific logic was handled by adding rule names (e.g.,
 477 “dynamic_price_floor”) as library entries or composing existing blocks (e.g., dynamic
 478 + cost_function). We provide the detailed mapping of these papers to MIND constructs in
 479 Appendix H.

486 **5 ABLATION STUDIES**

488 To quantify the impact of our data curation process, we conduct ablation studies on progressively
489 less-filtered versions of our training dataset: (1) **Parse-Only**: filtered only for syntactic correctness;
490 (2) **No Execute Check**: filtered through compilation (Parse → Validator → Compiler) but without
491 execution-time validation; (3) **No LLM Check**: passes full 4-stage validation (Parse → Validator
492 → Compiler → Execute) but lacks the description-DSL consistency verification. We train separate
493 MarketCopilot models on each dataset variant using identical architectures, training budgets, and
494 hyperparameters.

495
496
497 Table 3: Ablation study results demonstrating the impact of each data curation stage. All models
498 trained with identical architectures and budgets.

500 Model	501 Parse Success (%)	501 Validation + Compilation Success (%)	501 IR Equivalence (%)	501 ΔIR vs Previous (pp)
502 Parse-Only	503 98.33	503 81.33	503 66.33	503 –
504 w/o Execute Check	504 98.67	504 92.33	504 68.67	504 +2.34
505 w/o LLM Check	505 100.00	505 99.33	505 72.33	505 +3.66
506 Full Pipeline	507 100.00	508 100.00	509 96.33	510 +24.00

506 Table 3 demonstrates that each curation stage contributes significantly to final performance. While
507 parse success remains uniformly high (>98%) across all variants—indicating that learning basic
508 DSL syntax is straightforward—the gaps emerge in semantic correctness. Validation and compilation
509 success improves from 81.33% to 100% as filtering stages are added, with the execution check
510 contributing 7.00 percentage points and validator checks contributing 11.00 percentage points from
511 the Parse-Only baseline.

512 Most critically, IR equivalence shows dramatic improvement: from 66.33% (Parse-Only) to 96.33%
513 (Full Pipeline), a total gain of 30.00 percentage points. **This 30-point gain is driven by the IR-level**
514 **validation and description-DSL alignment, not by a larger LLM.** The description-DSL consistency
515 check alone contributes 24.00 percentage points (72.33% to 96.33%), highlighting that alignment
516 between natural language and formal specifications is crucial for semantic correctness. Without
517 this final verification, models generate syntactically valid but semantically incorrect programs—
518 they learn surface patterns rather than the underlying mapping between economic concepts and their
519 formal representations.

520 These results validate our design choice to prioritize data quality over quantity. Training on carefully
521 curated examples produces models that understand the semantic correspondence between natural
522 language descriptions and market mechanisms, rather than merely mimicking syntactic patterns.

524 **6 CONCLUSION AND FUTURE WORK**

527 We present MIND, a symbolic language and toolchain that bridges economic design and executable
528 implementation. Through a typed IR, phased validation, and dual-backend execution, MIND re-
529 duces specification complexity by 79% and enables an AI Copilot to achieve 96.33% semantic
530 correctness. **A within-subjects user study confirms that practitioners modify mechanisms 4–10×**
531 **faster in MIND than in Python. We further demonstrate generalization by encoding 8 mechanisms**
532 **from recent ACM EC/SIGecom literature.** Designed to be front-end agnostic, MIND allows stronger
533 reasoning models to plug into its governance layer—comprising spec hashes, validator reports, and
534 audit logs—to ensure traceability and compliance.

535 Future work focuses on expanding scope and integration. We plan to support combinatorial bidding
536 languages, iterative auctions, and stochastic Bayesian games, while scaling to millions of partic-
537 ipants. Crucially, we aim to bridge MIND with external ecosystems: this includes handling schema
538 drift in messy real-world data and integrating with formal verification engines (Passmore et al.,
539 2020) to mathematically prove economic properties. With these efforts, MIND provides a practical
foundation for exploring, auditing, and deploying market designs.

REPRODUCIBILITY STATEMENT

Our work is committed to the principles of open and reproducible research. To this end, all code, datasets, and experimental configurations will be made publicly available upon acceptance of this paper.

REFERENCES

Huda Alawad, Ramesh Panta, Minhaz Zibran, and Mohammad Amin Al Islam. An empirical study of the relationships between code readability and software complexity. In *2019 IEEE 27th International Conference on Program Comprehension (ICPC)*, pp. 108–118. IEEE, 2019. doi: 10.1109/ICPC.2019.00023.

Jalyawat Anunrojwong, Santiago Balseiro, and Omar Besbes. On the robustness of second-price auctions in prior-independent mechanism design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 87–88, 2022.

Ali Aouad and Will Ma. A nonparametric framework for online stochastic matching with correlated arrivals. *Management Science*, 69(10):6024–6046, 2023.

Luca Ardito, Rosella Coppola, Lorenzo Barbato, and Daniele Verga. A tool-based perspective on software code maintainability metrics: A systematic literature review. *Software: Practice and Experience*, 50(12):2203–2230, 2020. doi: 10.1002/spe.2876.

Lawrence M Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, 2004.

Martino Banchio and Andrzej Skrzypacz. Artificial intelligence and auction design. *Management Science*, 68(12):8749–8768, 2022.

N Bertram, A Levinson, and J Hsu. Cutting the cake: a language for fair division. [corr/abs/2304.04642](https://arxiv.org/abs/2304.04642) (2023), 2023.

Andrei Borshchev and Alexei Filippov. From system dynamics and discrete event to practical agent based modeling: reasons, techniques, tools. 2004.

Holger Stadel Borum and Christoph Seidl. Survey of established practices in the life cycle of domain-specific languages. In *Proceedings of the 25th International Conference on Model Driven Engineering Languages and Systems*, pp. 266–277, 2022.

MA Bouaicha, G Destefanis, T Montanaro, N Lasla, and L Patrono. Shill bidding prevention in decentralized auctions using smart contracts. *Information Sciences*, pp. 122374, 2025.

David Byrd, Maria Hybinette, and Tucker Hybinette Balch. Abides: Towards high-fidelity market simulation for ai research. *arXiv preprint arXiv:1904.12066*, 2019. URL <https://arxiv.org/abs/1904.12066>.

Rodrigo N. Calheiros, Rajiv Ranjan, Anton Beloglazov, César A. F. De Rose, and Rajkumar Buyya. Cloudsim: A toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Software: Practice and Experience*, 41(1):23–50, 2011. doi: 10.1002/spe.995. URL <https://onlinelibrary.wiley.com/doi/10.1002/spe.995>.

Marco B Caminati, Manfred Kerber, Christoph Lange, and Colin Rowat. Certified vickrey auctions. In *CADE*, 2015.

Christopher Cashore, Peter Frazier, and Eva Tardos. Dynamic pricing provides robust equilibria in stochastic ridesharing networks. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 257–258, 2022.

Dave De Jonge and Dongmo Zhang. Gdl as a unifying domain description language for declarative automated negotiation. *Autonomous Agents and Multi-Agent Systems*, 35(1):13, 2021.

594 Greg d’Eon, Neil Newman, and Kevin Leyton-Brown. Understanding iterative combinatorial auc-
595 tion designs via multi-agent reinforcement learning. In *Proceedings of the 25th ACM Conference*
596 *on Economics and Computation*, pp. 1102–1130, 2024.

597

598 Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex
599 optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

600 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
601 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
602 *arXiv e-prints*, pp. arXiv–2407, 2024a.

603

604 Avinava Dubey, Zhe Feng, Rahul Kidambi, Aranyak Mehta, and Di Wang. Auctions with llm
605 summaries. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and*
606 *Data Mining*, pp. 713–722, 2024b.

607

608 Paul Duetting, Vahab Mirrokni, Renato Paes Leme, Haifeng Xu, and Song Zuo. Mechanism design
609 for large language models. In *Proceedings of the ACM Web Conference 2024*, pp. 144–155, 2024.

610

611 Michal Feldman, Simon Goke, Brendan Lucier, and Renato Paes Leme. Bayesian and randomized
612 clock auctions. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp.
613 625–626, 2022.

614

615 David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American*
616 *Mathematical Monthly*, 69(1):9–15, 1962.

617

618 Simon Goke, Brendan Lucier, Renato Paes Leme, and Q Wang. Bidders’ responses to auction format
619 change: Evidence from a field experiment in display ad auctions. In *Proceedings of the 23rd ACM*
620 *Conference on Economics and Computation*, 2022.

621

622 Kachina Gui, Peijin Nair, Jian Niu, et al. Auction throttling and causal inference in unified optimiza-
623 tions of bidder pacing and publisher yield. In *Proceedings of the ACM Web Conference 2022*, pp.
624 123–133, 2022.

625

626 Samaneh Hoseindoost, Afsaneh Fatemi, and Bahman Zamani. An executable domain-specific mod-
627 eling language for simulating organizational auction-based coordination strategies for crisis re-
628 sponse. *Simulation Modelling Practice and Theory*, 131:102880, 2024.

629

630 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
631 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

632

633 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
634 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv:2410.21276*,
635 2024.

636

637 Hyejin Ko and Kamesh Munagala. Optimal price discrimination with a public budget. In *Proceed-
638 ings of the 23rd ACM Conference on Economics and Computation*, pp. 943–944, 2022.

639

640 Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, Jacques Pienaar, River
641 Riddle, Tatiana Shpeisman, Nicolas Vasilache, and Oleksandr Zinenko. Mir: Scaling compiler
642 infrastructure for domain specific computation. In *2021 IEEE/ACM International Symposium on*
643 *Code Generation and Optimization (CGO)*, pp. 2–14. IEEE, 2021.

644

645 Paul R. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press,
646 Cambridge, 2004. doi: 10.1017/CBO9780511813825. URL <https://www.cambridge.org/core/books/putting-auction-theory-to-work/63FA2A2D332E9E3A3238B6D5B650F2A5>.

647

648 Paul R. Milgrom. Auction research evolving: Theorems and market designs. *American Eco-
649 nomic Review*, 111(5):1383–1405, 2021. doi: 10.1257/aer.111.5.1383. URL <https://www.aeaweb.org/articles?id=10.1257/aer.111.5.1383>.

650

651 Andreea Molnar and Simona Motogna. Longitudinal evaluation of open-source software main-
652 tainability. *arXiv preprint arXiv:2003.00447*, 2020. URL <https://arxiv.org/abs/2003.00447>.

648 Noam Nisan. Bidding languages for combinatorial auctions. In *Combinatorial Auctions*, pp. 215–
649 231. MIT Press, 2006.

650

651 Curtis Northcutt, Lu Jiang, and Isaac Chuang. Confident learning: Estimating uncertainty in dataset
652 labels. *Journal of Artificial Intelligence Research*, 70:1373–1411, 2021.

653

654 Grant Olney Passmore, Denis Ignatovich, Kae Aitken, Konstantin Anisimov, M Bray, et al. The
655 imandra automated reasoning engine. *arXiv preprint arXiv:2010.03215*, 2020.

656

657 Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré.
658 Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the VLDB en-*
659 *dowment. International conference on very large data bases*, volume 11, pp. 269, 2017.

660

661 Alvin E. Roth. Marketplaces, markets, and market design. *American Economic Review*, 108
662 (7):1609–1658, 2018. doi: 10.1257/aer.108.7.1609. URL <https://www.aeaweb.org/articles?id=10.1257/aer.108.7.1609>.

663

664 Anand Shah, Kehang Zhu, Yanchen Jiang, Jeffrey G Wang, Arif K Dayi, John J Horton, and David C
665 Parkes. Learning from synthetic labs: Language models as auction participants. *arXiv preprint*
666 *arXiv:2507.09083*, 2025.

667

668 Yu-Zhe Shi, Haofei Hou, Zhangqian Bi, Fanxu Meng, Xiang Wei, Lecheng Ruan, and Qining Wang.
669 Autodsl: Automated domain-specific language design for structural representation of procedures
670 with constraints. *arXiv preprint arXiv:2406.12324*, 2024.

671

672 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
673 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing
674 Systems*, 36:8634–8652, 2023.

675

676 Jie Sun, Tianyu Zhang, Houcheng Jiang, Kexin Huang, Chi Luo, Junkang Wu, Jiancan Wu,
677 An Zhang, and Xiang Wang. Large language models empower personalized valuation in auc-
678 tion. *arXiv preprint arXiv:2410.15817*, 2024.

679

680 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
681 Chang Gao, Chengan Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*
682 *arXiv:2505.09388*, 2025.

683

684 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
685 React: Synergizing reasoning and acting in language models. In *International Conference on
686 Learning Representations (ICLR)*, 2023.

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 USE OF LLMs IN OUR WORK 703

704 We used large language models (LLMs) in four ways: (i) manuscript polishing—to improve grammar,
705 clarity, and flow without altering substantive claims; (ii) literature triage—to surface potentially
706 relevant papers; (iii) data creation—to synthesize a portion of our NL→DSL pairs (see Sec. 3.4);
707 and (iv) *prompt design*—to iterate on task instructions and few-shot exemplars. All LLM outputs
708 affecting results were reviewed by authors for accuracy, and dataset items were validated with our
709 parser/validator pipeline and spot-audited by humans.

710 711 A MARKET INTERPRETATION DSL COMPONENTS 712

713 The core building blocks and their overall structure are:
714

```
715
716 market {
717     global_settings { ... }
718     auction { ... } // repeatable
719     stage(name="...") { ... } // optional, repeatable
720     matching { ... } // optional
721     constraints { ... } // optional
722     objectives { ... } // optional
723     dynamic { ... } // optional
724 }
```

725 Block Descriptions 726

- 727 • **market**: Top-level container for a complete market specification.
- 728 • **global_settings** (optional): Global parameters (e.g., units, supply, defaults).
- 729 • **auction** (repeatable): Auction mechanism definition (type, participants, goods, valuations, rules,
730 distributions).
- 731 • **stage** (optional, repeatable): Multi-stage orchestration with re-auctioning and discounting logic.
- 732 • **matching** (optional): Matching market (type, participants, compatibility, rule).
- 733 • **constraints** (optional): Feasibility/policy conditions; simple or parameterized forms.
- 734 • **objectives** (optional): Optimization goals used by solver backends.
- 735 • **dynamic** (optional): [Settings for time-varying parameters, multi-period loops, and discounting.](#)

736 737 A.1 SUPPORTED MECHANISM SCOPE 738

739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756
757
758

Table 4: MIND Scope of Expressiveness. We categorize features into currently supported (In-Scope) and those reserved for future work (Out-of-Scope).

759 760 761 762 763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799 800 801 802 803 804 805 806 807 808 809	Feature Category	Supported (In-Scope)	Out-of-Scope
	Auction Formats	<ul style="list-style-type: none"> Standard: First-Price, Second-Price, Uniform-Price, Pay-as-Bid. Combinatorial: Additive valuations with logical constraints (AND/OR, XOR via constraints). 	<ul style="list-style-type: none"> Complex Combinatorial: General XOR bidding languages without constraints. Iterative: Ascending clock auctions (except via custom dynamic loops).
	Matching Logic	<ul style="list-style-type: none"> Bipartite: One-to-one, one-to-many. Stable Matching: Deferred Acceptance logic. Compatibility: Explicit allow/deny graphs. 	<ul style="list-style-type: none"> General Graph: Non-bipartite matching with long exchange cycles (e.g., kidney chains > 2).
	Dynamics & Stochasticity	<ul style="list-style-type: none"> Distributions: Poisson, Uniform, Custom discrete. Randomness: Drawing random values for valuations/supply. Time Loops: Multi-period execution with discounting. Policy Functions: Dynamic reserves/costs based on expressions. 	<ul style="list-style-type: none"> Stochastic Optimization: Solving <i>for</i> optimal policies under uncertainty (MIND simulates policies, it doesn't <i>solve</i> them). Equilibrium Finding: Computing Nash equilibria.
	Constraints	<ul style="list-style-type: none"> Linear: Budget, supply, capacity, fairness quotas. Logical: Mutual exclusivity, package constraints. Incentive: IR, IC (as linear constraints). 	<ul style="list-style-type: none"> Non-Convex: Integer constraints not mappable to MIP. Black-box: Constraints defined by external code.
	Objectives	<ul style="list-style-type: none"> Convex: Maximize Revenue, Welfare, Matches, or custom convex functions. 	<ul style="list-style-type: none"> Non-Convex: Deep neural network objectives. Game-Theoretic: “Maximize stability” (unless expressed as a matching rule).

B FORMAL GRAMMAR (EBNF)

```

program           : "market" "(" market_block* ")" ;
market_block     : global_settings_block
                  | auction_block
                  | stage_block
                  | matching_block
                  | constraints_block
                  | objectives_block ;
global_settings_block
: "global_settings" "(" gs_element* ")" ;
gs_element        : currency_decl
                  | supply_decl
                  | reserve_price_decl ;
auction_block     : "auction" "(" auction_element* ")" ;
auction_element   : auction_type_decl
                  | participants_decl
                  | goods_decl
                  | valuations_decl
                  | allocation_rule_decl
                  | payment_rule_decl ;

```

```

810  stage_block      : "stage" "(" "name" "=" string_literal ")" "
811      | " stage_element* ")" ;
812  stage_element    : auction_block
813      | reauction_decl ;
814  reauction_decl   : "reauction" "(" "unsold_goods" "=" string_literal ","
815      | "auction_type" "=" string_literal ")" ;
816
817  matching_block   : "matching" "{" matching_element* "}" ;
818  matching_element  : matching_type_decl
819      | participants_decl
820      | compatibility_graph_decl
821      | matching_rule_decl ;
822
823  constraints_block: "constraints" "{" constraint_entry_list "}" ;
824  objectives_block  : "objectives" "(" objective_entry_list ")" ;
825
826  participants_decl : "participants" "(" string_literal_list ")" ;
827  goods_decl        : "goods" "(" string_literal_list ")" ;
828  valuations_decl   : "valuations" "(" valuation_entry_list ")" ;
829
830  string_literal_list: "[" (string_literal (," string_literal)*)? "]" ;
831  valuation_entry_list:
832      : "{" (valuation_entry (," valuation_entry)*)? "}" ;
833  valuation_entry   : string_literal ":" "{" good_value (," good_value)* "}" ;
834  good_value         : string_literal ":" number ;
835
836  auction_type_decl : "type"      "(" string_literal ")" ;
837  allocation_rule_decl:
838      : "allocation_rule" "(" string_literal ")" ;
839  payment_rule_decl  : "payment_rule" "(" string_literal ")" ;
840
841  matching_type_decl : "type"      "(" string_literal ")" ;
842  matching_rule_decl  : "matching_rule" "(" string_literal ")" ;
843  compatibility_graph_decl:
844      : "compatibility_graph" "(" compatibility_entry_list ")" ;
845  compatibility_entry_list:
846      : "{" (compatibility_entry (," compatibility_entry)*)? "}" ;
847  compatibility_entry: string_literal ":" "[" (string_literal (," string_literal)*)? "]" ;
848
849
850  constraint_entry_list:
851      : (constraint_param_entry | string_literal)
852      | (constraint_param_entry | string_literal)* ;
853  constraint_param_entry:
854      : identifier "(" (parameter_assignment
855      | (," parameter_assignment)*)? ")" ;
856  parameter_assignment:
857      : identifier "=" value ;
858
859  objective_entry_list:
860      : (string_literal (," string_literal)*)? ;
861
862  string_literal     : ESCAPED_STRING ;
863  identifier         : /[A-Za-z_][A-Za-z0-9_]/ ;
864  number              : SIGNED_NUMBER ;
865  value               : number | string_literal | boolean ;
866  boolean             : "true" | "false" ;
867
868
869

```

C VALIDATION

What is verified

- Names and References: unique goods/participants; auctions reference declared participants/goods.
- Valuations Consistency: keys match auction participants; goods in valuations are declared; sparse entries → warnings.
- Rules Recognition: auction types, allocation/payment rules recognized or mapped from common aliases.
- Stage/Matching (if present): global settings sanity; stage naming; reauction fields; matching type/rule; graph nodes exist; symmetry warnings.
- Constraints/Objectives (if present): types recognized; basic parameter sanity; objective conflict warnings.

864

Simple Validation Algorithm

865

```

866 Input           : MarketProgram
867 Output          : ValidationReport (errors, warnings, suggestions); IR may be autofixed
868
869 1) Basic field checks (hard errors)
870   - Good.name not empty; reserve_price >= 0
871   - Bidder.name not empty; budget >= 0
872   - Auction.auction_type not empty
873   - Assignment fields not empty; bid_price >= 0
874
875 2) Core checks (always)
876   - Unique names; auctions reference existing participants/goods
877   - valuations match auction participants/goods; sparse -> WARN, mismatches -> ERROR
878   - auction_type, allocation_rule, payment_rule:
879     - map known aliases
880     - unknown -> ERROR; some types partially implemented -> WARN
881
882 3) Stage/Matching checks (only if present)
883   - Global settings: supply/reserve defaults; negatives -> ERROR, missing -> WARN
884   - Stages: unique, named, each has an auction
885   - Reauction: needs unsold_goods and auction_type; validate type; loose goods check
886   - Matching: normalize matching_type; unknown -> WARN + default to bipartite
887     - participants non-empty, no duplicates
888     - compatibility_graph nodes exist; symmetry missing -> WARN
889     - matching_rule: missing -> default stable_matching (WARN); unknown -> ERROR
890
891 4) Advanced checks (only if constraints/objectives present)
892   - Constraints: type recognized; params sane (e.g., budgets >= 0, caps >= 0)
893   - Objectives: normalize; unknown -> ERROR; conflicting goals -> WARN
894
895 5) Autofix safe defaults
896   - Missing/unknown allocation_rule -> highest_bid
897   - Missing/unknown payment_rule -> second_price
898   - Global supply missing/invalid -> 1
899   - Missing reserve_price (global/good) -> 0.0
900   - All fixes logged as suggestions
901
902 6) Return report; program contains applied defaults where safe
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

```

C.1 GOVERNANCE ARTIFACTS

To support auditing and compliance, MIND treats the compilation process as a governance event. A complete **Governance Artifact** is a versioned bundle containing four components that allow third parties to verify outcomes:

1. **Readable Specification:** The human-legible MIND DSL source code defining the mechanism rules.
2. **Typed IR Snapshot:** The serialized Intermediate Representation (AST) used by the compiler, serving as the single source of truth for semantics.
3. **Validator Report:** A structured log of all checks passed (parsing, typing, economic consistency) and any auto-fixes applied.
4. **Execution Integrity:** Cryptographic hashes of the spec and IR, logged alongside backend simulation outputs (allocations, payments) to ensure reproducibility.

Together, these components decouple the *intent* (DSL) from the *execution* (Backend), providing a transparent chain of custody for market mechanisms.

D BACKEND SELECTION

Heuristic (implemented in **MarketCompiler**)

- Pure matching or simple Phase-1 auctions → Pandas (NetworkX for matching).
- Multi-stage without optimization features → Pandas + Prefect orchestration.
- Combinatorial auctions or constraints/objectives → CVXPY optimization.

918 Mapping

919

920 IR features	921 Selected backend
922 auction only, valuations, simple rules	Pandas
923 matching (bipartite/stable)	Pandas + NetworkX
924 stage flow (no constraints/objectives)	Pandas + Prefect
925 constraints/objectives present	CVXPY
926 combinatorial auction	CVXPY

927 E DATASET GENERATION PIPELINE

928 Overview

- 931 1. Sample a use case (914 total) and random market settings; assemble 4-shot prompt + grammar
(markdown + EBNF).
- 933 2. Generate (GPT-4o-mini) brief description + DSL.
- 934 3. Filter with 4-stage validation (Parser → Validator → Compiler → Execute); keep only programs
935 that pass all stages.
- 936 4. Enhance description (GPT-4o-mini) by extracting all facts from DSL; replace description text
937 only.
- 938 5. Consistency check (GPT-4o): “YES/NO” whether description matches DSL; keep YES, drop
939 NO.
- 940 6. **Manual Inspection & De-duplication:** Final manual review and strict de-duplication to form the
941 final dataset.

943 E.1 DATA FILTERING STATISTICS

945 To guarantee dataset quality, we tracked the number of samples retained at each stage. Table 5
946 details the rigorous filtering process.

949 Table 5: Dataset Filtering Pipeline Statistics. The high drop rate at the Consistency Check stage
950 ensures that only accurately described mechanisms are retained.

951 Pipeline Stage	952 Input Count	953 Dropped	954 Remaining
952 Raw Generation (GPT-4o-mini)	953 20,108	954 –	955 20,108
953 Stage 1 & 2: Parse + Validator	954 20,108	955 551	956 19,557
954 Stage 3: Compiler	955 19,557	956 302	957 19,255
955 Stage 4: Execute	956 19,255	957 186	958 19,069
956 Stage 5: LLM Consistency Check	957 19,069	958 6,340	959 12,729
960 Validated Pool			961 12,729

962 Prompt for data generation

963 You are an expert Market Mechanism DSL generator. I will provide you with:

- 964 – The grammar of the MarketMechanismDSL,
- 965 – A few example DSL programs,
- 966 – A target use case,
- 967 – And a set of market settings.

968 Your task is to generate a complete MarketMechanismDSL program that
969 fits the given scenario and settings. Strictly follow the provided
970 grammar and take inspiration from the examples. Use the canonical
971 constructs: participants([...]), goods([...]), valuations({ ... }),
972 and valid allocation_rule/payment_rule names. Do not invent syntax
973 not present in the grammar.

974 Please output your response in EXACTLY the following format and nothing else:

```

972
973     Description:
974
975     ```markdown
976     1-3 sentences written from the user's perspective describing the market
977     ```

978     DSL code:
979
980     ```dsl
981     <your complete MarketMechanismDSL program here>
982     ```

983     Inputs:
984
985     DSL Grammar:
986     ```markdown
987     {grammar}
988     ```

989     DSL Program Examples:
990     ```dsl
991     {examples}
992     ```

993     Use Case and Settings:
994     - Domain: {domain}
995     - Scenario: {scenario}
996     - Settings: {settings}
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

```

Prompt for description completion (from DSL)

```

996     You are a precise technical writer. Given a MarketMechanismDSL program,
997     write a COMPLETE, human-readable task description that includes ALL facts
998     present in the DSL (participants, goods, valuations, auction/matching type,
999     allocation/payment/matching rules, key settings). Do NOT hallucinate new
1000    entities or numbers. Use clear, concise prose (4-8 sentences).
1001
1002    Output EXACTLY in this format:
1003
1004    Description:
1005
1006    ```markdown
1007    <concise but complete description, entirely derived from the DSL>
1008    ```

1009    Inputs:
1010
1011    DSL Program:
1012    ```dsl
1013    {dsl_code}
1014    ```

1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

```

F COMPLETENESS AGENT WORKFLOW

Node-by-node extraction Given a natural-language description, the agent queries nodes in order and marks each as *Enough/Not Enough*:

1. Auction (type, participants, goods, valuations, rules).
2. Global settings (currency, supply, reserve_price).
3. Stages (name, auction, optional reauction).
4. Matching (type, participants, compatibility_graph, rule).

1026 5. Constraints (simple/parameterized).
 1027 6. Objectives (list).

1029 If Not Enough, it asks for the missing facts before proceeding. Finally it merges fragments into a
 1030 normalized schema and renders minimal DSL.
 1031

1032 **G END-TO-END EXAMPLE (SIMPLE AUCTION)**

1033 **G.1 NATURAL LANGUAGE INPUT**

1036 **Step 1: Natural Language**

1038 "Create a second-price auction for three participants competing for an advertising slot.
 1039 Alice values the slot at \$100, Bob at \$80,
 1040 and Carol at \$90. Use highest bidder wins allocation."

1042 **G.2 COPILOT GENERATION (NL → DSL)**

1044 **Step 2: Generated DSL**

```
1046 market {
1047   auction {
1048     type("second_price")
1049     participants(["Alice", "Bob", "Carol"])
1050     goods(["AdSlot"])
1051     valuations({
1052       "Alice": {"AdSlot": 100},
1053       "Bob": {"AdSlot": 80},
1054       "Carol": {"AdSlot": 90}
1055     })
1056     allocation_rule("highest_bid")
1057     payment_rule("second_price")
1058   }
1059 }
```

1058 **G.3 PARSER (DSL → PARSE TREE)**

1060 **Step 3: Lark Parse Tree (excerpt)**

```
1062 program
1063   statement
1064     auction_block
1065       auction_element
1066         auction_type_decl
1067           string_literal: "second_price"
1068         auction_element
1069           participants_decl
1070             string_literal_list
1071               string_literal: "Alice"
1072               string_literal: "Bob"
1073               string_literal: "Carol"
1074 ...
1075
```

1074 **G.4 TRANSFORMER (PARSE TREE → IR)**

1075 **Step 4: Intermediate Representation (MarketProgram)**

```
1076 MarketProgram(
1077   goods=[Good(name="AdSlot")],
1078   bidders=[Bidder(name="Alice"), Bidder(name="Bob"), Bidder(name="Carol")],
1079   auctions=[Auction(
1080     auction_type="second_price",
```

```
1080
1081     participants=["Alice","Bob","Carol"],
1082     goods=["AdSlot"],
1083     valuations={"Alice":{"AdSlot":100}, "Bob": {"AdSlot":80}, "Carol": {"AdSlot":90}},
1084     allocation_rule="highest_bid",
1085     payment_rule="second_price"
1086   ],
1087   assignments=[], global_settings=None, stages=[],
1088   matching=None, constraints=[], objectives=[]
1089 )
```

G.5 VALIDATION (IR SEMANTIC CHECK)

Step 5: Validation

```
1093     program.validate_comprehensive() successful
1094     - All participants referenced in valuations
1095     - All goods have valid valuations
          - Auction type and payment rule compatible
          - Unique participant and good names verified
```

G.6 COMPILATION AND BACKEND SELECTION

Step 6: Compiler Analysis

```
1102 Analyzing DSL features...
1103 - Single auction block detected
1104 - No constraints or objectives
1105 - No multi-stage features
1106 - No matching markets
1107
1108 Backend Selected: Pandas
1109 Reasoning: Simple auction, simulation-based approach sufficient
```

G.7 EXECUTION (GENERATED CODE EXCERPT)

Step 7: Pandas Backend Code (excerpt)

```

1113 # Allocation: highest_bid rule
1114 allocations = {}
1115 for good in goods:
1116     good_bids = bids_df[bids_df['good'] == good].copy()
1117     if not good_bids.empty:
1118         winner = good_bids.loc[good_bids['valuation'].idxmax()]
1119         allocations[good] = winner['participant']
1120
1121 # Payment: second_price rule
1122 payments = {}
1123 for good, winner in allocations.items():
1124     good_bids = bids_df[bids_df['good'] == good].copy()
1125     sorted_bids = good_bids.sort_values('valuation', ascending=False)
1126     if len(sorted_bids) >= 2:
1127         second_highest = sorted_bids.iloc[1]['valuation']
1128         payments[winner] = payments.get(winner, 0) + second_highest
1129     else:
1130         payments[winner] = payments.get(winner, 0) + sorted_bids.iloc[0]['valuation']

```

G.8 RESULTS

Step 8: Console Output

```
1131      === AUCTION RESULTS ===  
1132      Allocations:  
1133          AdSlot: Alice  
1133          Payments:
```

Alice: \$90
Artifacts:
- auction_results.csv

1138
1139
1140

H ACADEMIC MECHANISM CASE STUDIES

1141
1142

1143 Table 6 illustrates how MIND captures diverse mechanisms from recent literature. We map the core
1144 theoretical components of each paper to specific MIND language constructs.
1145
1146

1147 Table 6: [Mapping theoretical components to MIND constructs for 8 real-world mechanisms.](#)
1148

Paper & Mechanism	Mapping to MIND Constructs	DSL Implementation (Snippet)
Banchio & Skrzypacz (2022) <i>Repeated First-Price Auction</i> Analyzes equilibrium in repeated auctions.	<ul style="list-style-type: none"> • Repeated Game → stage block with discount. • Equilibrium Condition → threshold w/ math expression. 	<pre>stage(name="RepeatedFPA") { auction { type("first_price") threshold(expr="(m-2) / (2m-3)") } discount(0.6) }</pre>
Gui et al. (2022) <i>Dynamic Second-Price Auction</i> Updates participation probability dynamically.	<ul style="list-style-type: none"> • Time Dynamics → dynamic block with periods. • Throttling Policy → cost_function with step(). 	<pre>dynamic { periods(288) } auction { type("second_price") cost_function(target="participation_prob", expr="step(r_t,[0.3,0.5],[0.3,0.5])") }</pre>
Feldman et al. (2022) <i>Uniform Price Auction</i> Verifies welfare bounds.	<ul style="list-style-type: none"> • Welfare Goal → objectives { "max_social_welfare" }. • Mechanism → Standard uniform_price type. 	<pre>auction { type("uniform_price") allocation_rule("uniform") } objectives { "maximize_social_welfare" }</pre>
Goke et al. (2022) <i>First-Price with Dynamic Floors</i> Real-time floor adjustment.	<ul style="list-style-type: none"> • Price Floor → threshold (reserve price). • Adjustment Logic → cost_function updating floor. 	<pre>auction { type("first_price") threshold(expr="dynamic_floor") cost_function(target="floor_policy", expr="update_rule") }</pre>
Cashore et al. (2022) <i>Stochastic Ridesharing (SSP)</i> Joint pricing and matching.	<ul style="list-style-type: none"> • Rider-Driver Matching → matching block (bipartite). • Pricing Stage → stage block with auction. 	<pre>matching { type("bipartite") matching_rule("maximize_matches") } stage(name="Pricing") { auction { type("second_price") } }</pre>
Anumrojwong et al. (2022) <i>Robust Second-Price Auction</i> Regret minimization with reserves.	<ul style="list-style-type: none"> • Regret Minimization → Implicitly via threshold tuning. • Reserve Price → Parameterized threshold. 	<pre>auction { type("second_price") threshold(expr="reserve_price") } constraints { "incentive_compatibility" }</pre>
Ko & Munagala (2022) <i>Optimal Randomized Auction</i> Revenue maximization under budget.	<ul style="list-style-type: none"> • Budget Limit → constraints { "budget_constraint" }. • Optimal Goal → objectives { "maximize_revenue" }. 	<pre>constraints { "budget_constraint" } objectives { "maximize_revenue" } stage(name="Stage1") { discount(0.6) }</pre>
Aouad & Ma (2023) <i>Online Stochastic Matching</i> Matching with correlated arrivals.	<ul style="list-style-type: none"> • Stochastic Demand → distribution (Custom/Poisson). • Expected Reward → cost_function expression. 	<pre>matching { type("bipartite") distribution(name="demand", kind="Custom", params={...}) cost_function(target="reward", expr="1-(1-rho) ^ 2") }</pre>

1188 **I** **USER STUDY DETAILS**
1189

1190 **I.1** **METHODOLOGY**
1191

1192 **Participants.** We recruited $N = 17$ participants with varying levels of programming and economic
1193 expertise. The breakdown of participants is as follows:

1194

- 1195 • **Education:** 13 undergraduate students, 4 graduate students (Master’s/PhD).
- 1196 • **Python Experience:** The cohort was technically proficient: 12 participants had > 2 years
1197 of experience, 2 had 1–2 years, and 3 had 0.5–1 year.
- 1198 • **Domain Knowledge:** 10 participants reported prior familiarity with auctions or matching
1199 markets, while 7 reported no prior domain knowledge.

1200 **Design.** The study followed a **within-subjects design**. Each participant performed modification
1201 tasks on the same three mechanisms in both **MIND** and **Python** environments. External tools (e.g.,
1202 IDEs, documentation, AI assistants) were allowed to mimic realistic workflows, provided that the
1203 total time spent was reported.

1204 **Tasks.** Participants started with working baseline code and were asked to implement specific policy
1205 modifications:

1206

- 1207 • **Task A (Pricing Rule):** Change the payment rule from second-price to first-price (pay-as-
1208 bid) while keeping allocation logic unchanged.
- 1209 • **Task B (Multi-Stage & Reserve):** Introduce a global reserve price (\$100) and a second
1210 “clearance” stage (reauction) for unsold goods.
- 1211 • **Task C (Compatibility):** Adjust the bipartite matching constraints (e.g., restrict specific
1212 Buyer-Tutor pairs and add a new agent) without changing the matching algorithm.

1213 **I.2** **RESULTS ANALYSIS**
1214

1215 Table 7 summarizes the performance metrics. MIND demonstrated a clear advantage in modification
1216 efficiency across all tasks.

1217 **Table 7: User Study Results ($N = 17$).** Comparison of completion time (median minutes) and
1218 correctness rate between MIND and Python workflows.

Task	Median Time (min)		Speed-up	Correctness Rate (%)	
	Python	MIND		Python	MIND
Task A (Pricing)	10.0	1.0	10.0 \times	88%	100%
Task B (Stages)	14.0	2.0	7.0 \times	76%	88%
Task C (Matching)	8.0	2.0	4.0 \times	82%	88%
Overall Average	10.7	1.7	6.3\times	82%	92%

1219 **Analysis.** The disparity was most pronounced in Task A and B. In Task A, MIND required only
1220 a keyword change, yielding a 10 \times speedup. In Task B, Python implementations often required
1221 significant refactoring of state management logic (14 min median), leading to higher time costs. In
1222 contrast, MIND allowed adding a stage in just 2 minutes. In Task C, direct manipulation of the
1223 compatibility_graph in MIND eliminated common index-error bugs observed in the Python
1224 group, improving both speed and correctness.