IMPACT: An Inference-Driven Modeling Framework for Cost-Effective Incentive Allocation in Service Operations

Yiheng An, Jia Li, Jeffrey D. Camm
School of Business, Wake Forest University

Liang Hu, Qinqin Zhuge, Bingxin Jia
Alibaba Group

KDD 2025 Workshop - Causal Inference and Machine Learning in Practice
August 4, 2025

Customer Satisfaction is Important

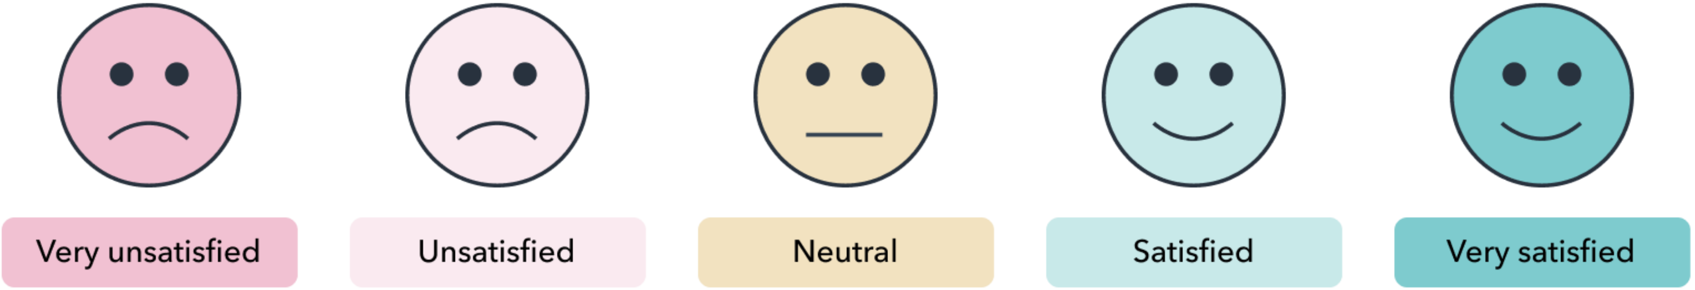

- Satisfied customers tend to stay with the company longer, spend more, and become brand advocates.
- Customer satisfaction also positively predicts a firm’s profit, stock returns, and other financial performance metrics.

Context: Taobao's Customer Service

Case: A "case" refers to an online service session between an agent and a customer where they communicate through text.

Case Difficulty: Difficulty level is labeled by a rule-based model ("Easy" vs "Difficult")

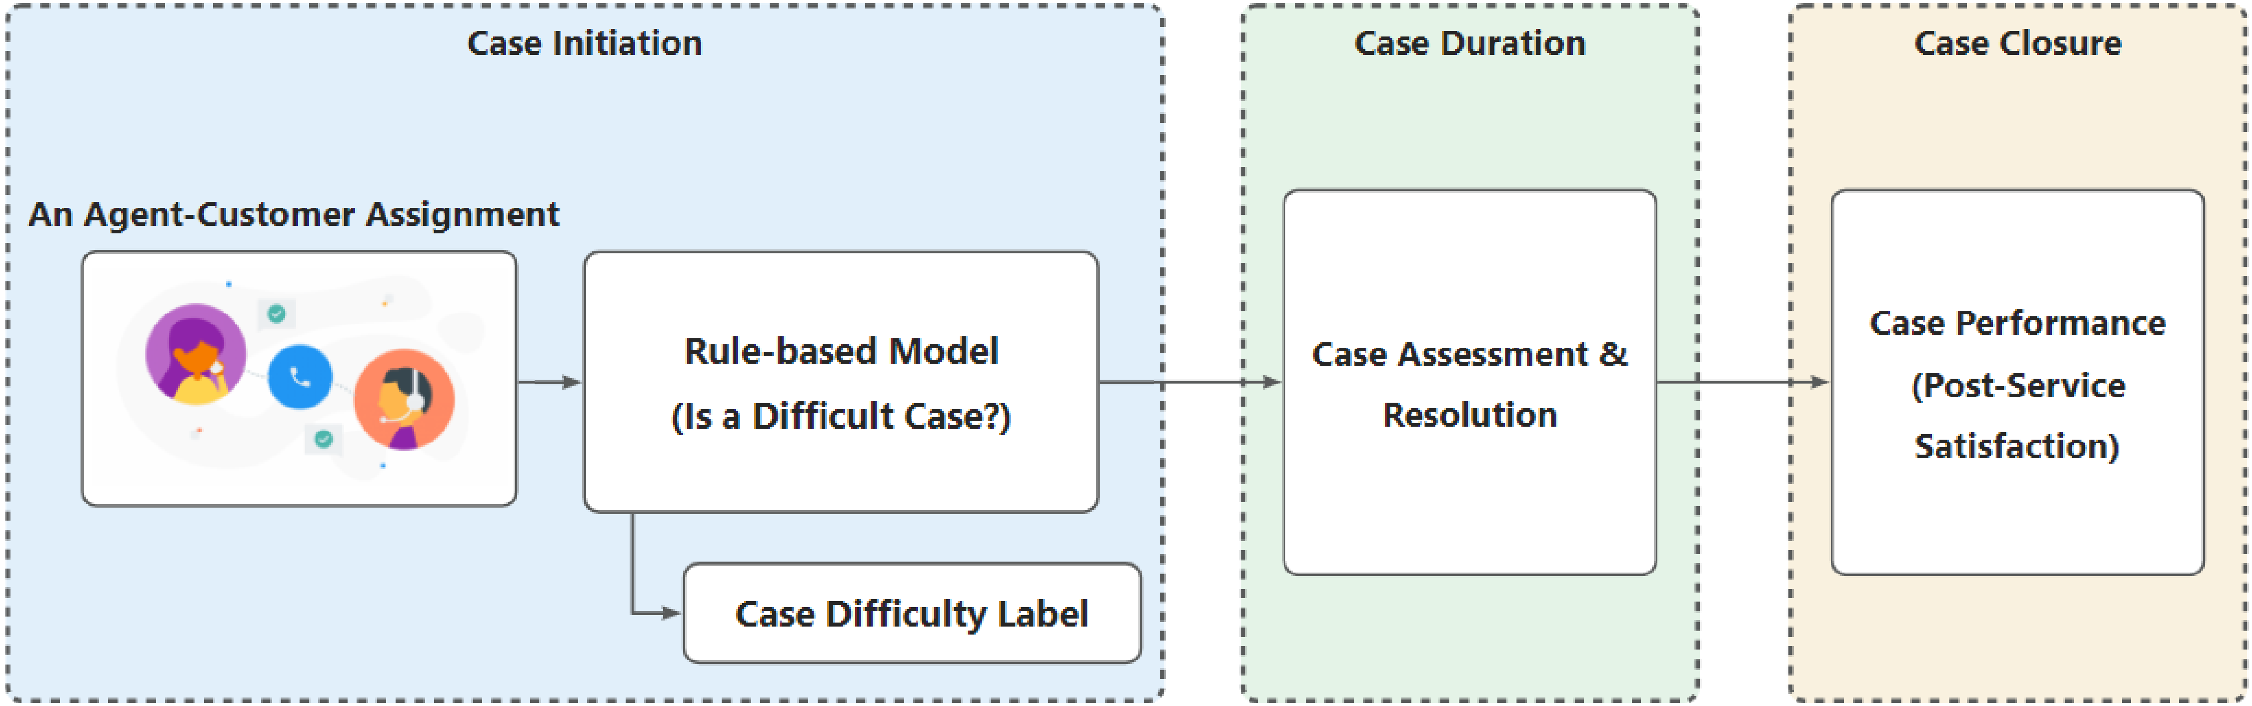

What We Did: A New Bonus System

- **Business Goal**: To enhance customer satisfaction by offering bonuses to human agents
- **Project Goal**: To design a new, performance-based bonus system for ***difficult cases***
- **Existing Approaches**: Largely rely on pre-defined rules, which typically lead to
  - ➢ limited adaptability to dynamic, context-specific information
  - ➢ inefficiency in measuring and optimizing ROI
  - ➢ lack of proactive capabilities at granular levels (e.g., case level)
- **Our Approach**:

  A context-aware, individualized, proactive bonus system that adapts to large-scale, high-frequency decision-making settings

The Bonus Offering Decision

The process of bonus notification and offering at Taobao customer service

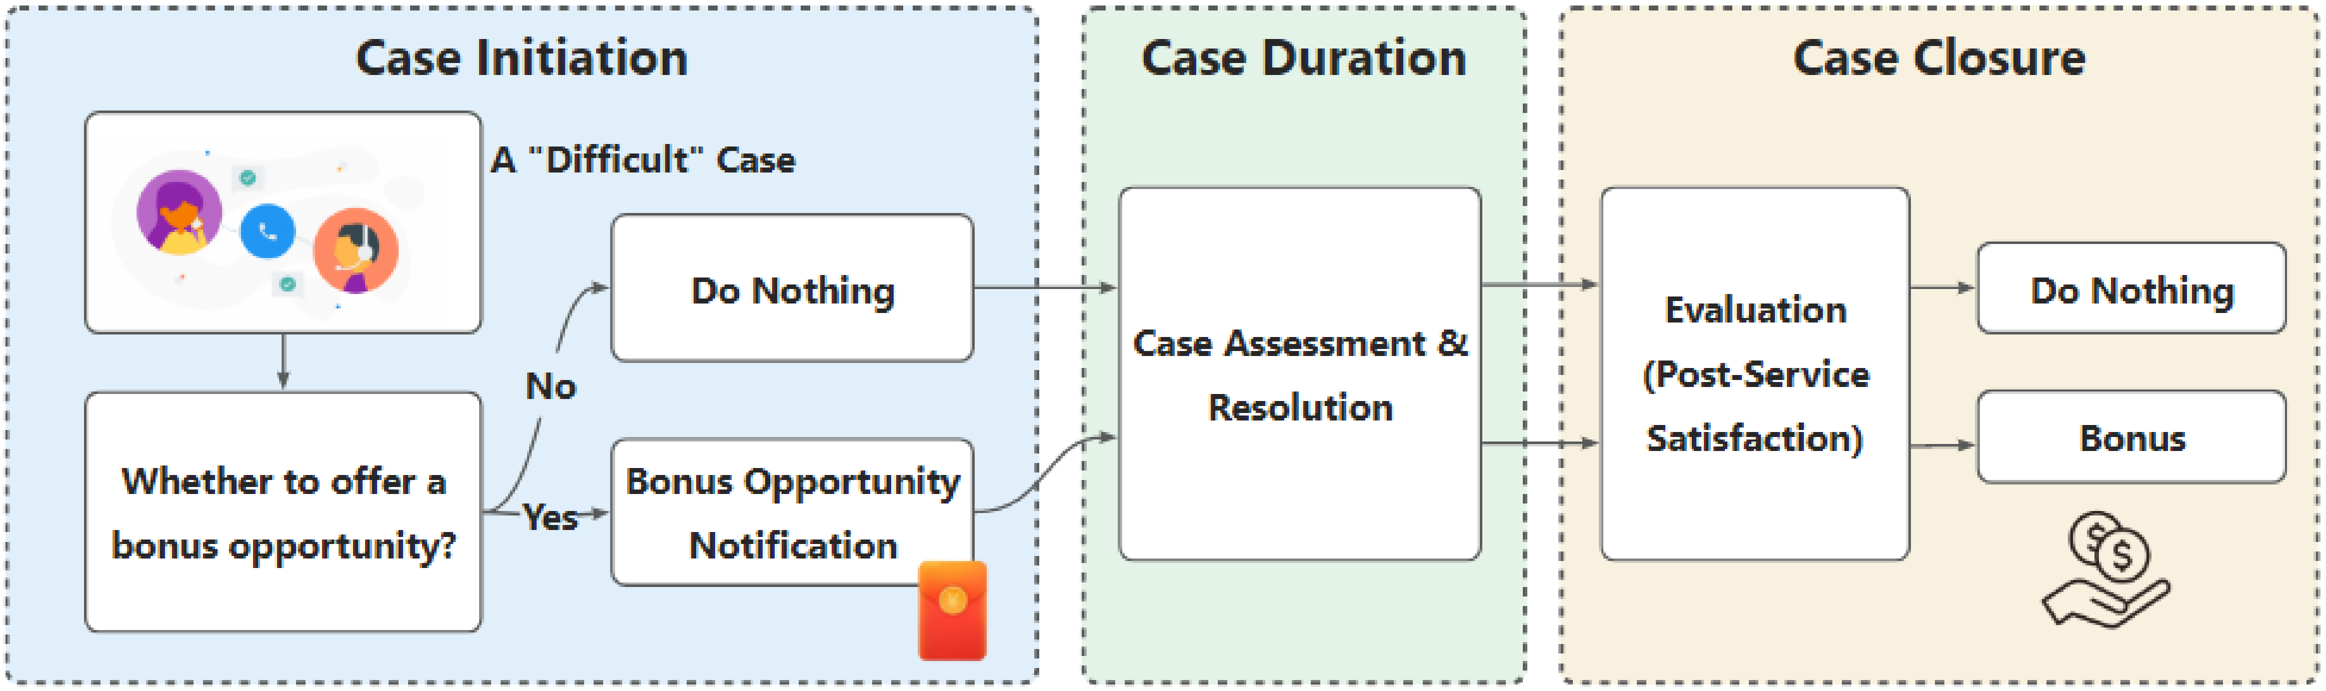

The Bonus Offering Decision

The process of bonus notification and offering at Taobao customer service

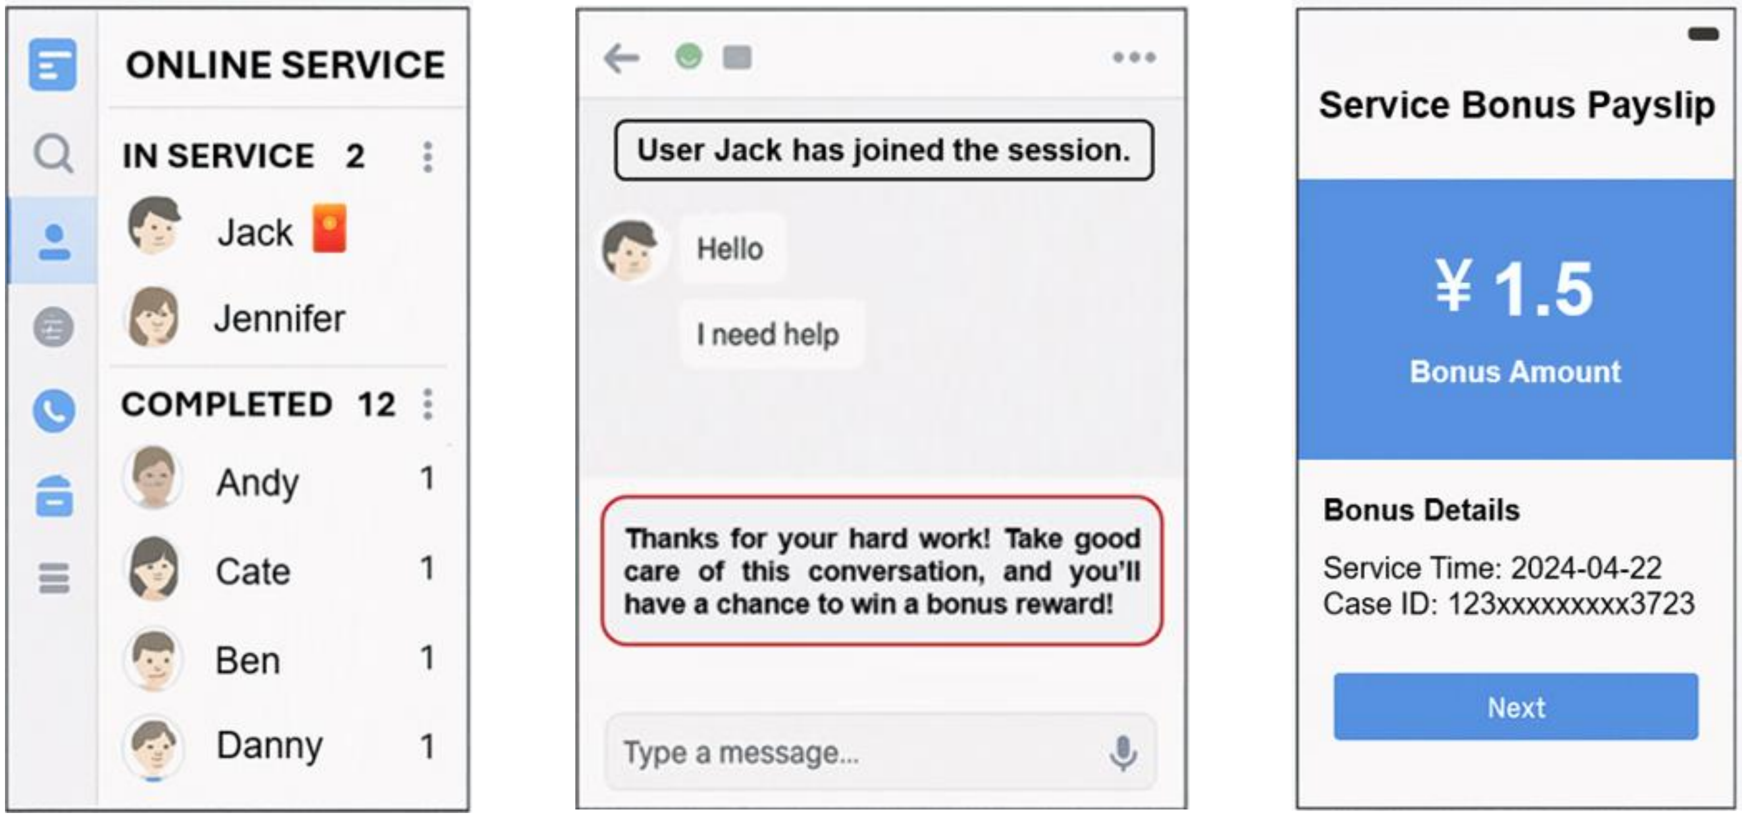

(a) Bonus Opportunity Notification

(b) Payment

The Bonus Targeting Problem

- **Objective**: Maximize the impact of bonuses on *customer satisfaction* for difficult cases
- **Decision**
  - Choice: Whether to offer a bonus opportunity to the agent given a *difficult* case
  - Timing: Case initiation phase, *before* interacting with the customer
  - Amount: Cash amount based on post-service customer satisfaction rating
- **Constraint**: Daily bonus budget

Call for

- Quantifying the case-specific effects of bonuses on customer satisfaction
- A decision-making mechanism that integrates context factors

IMPACT: System Overview

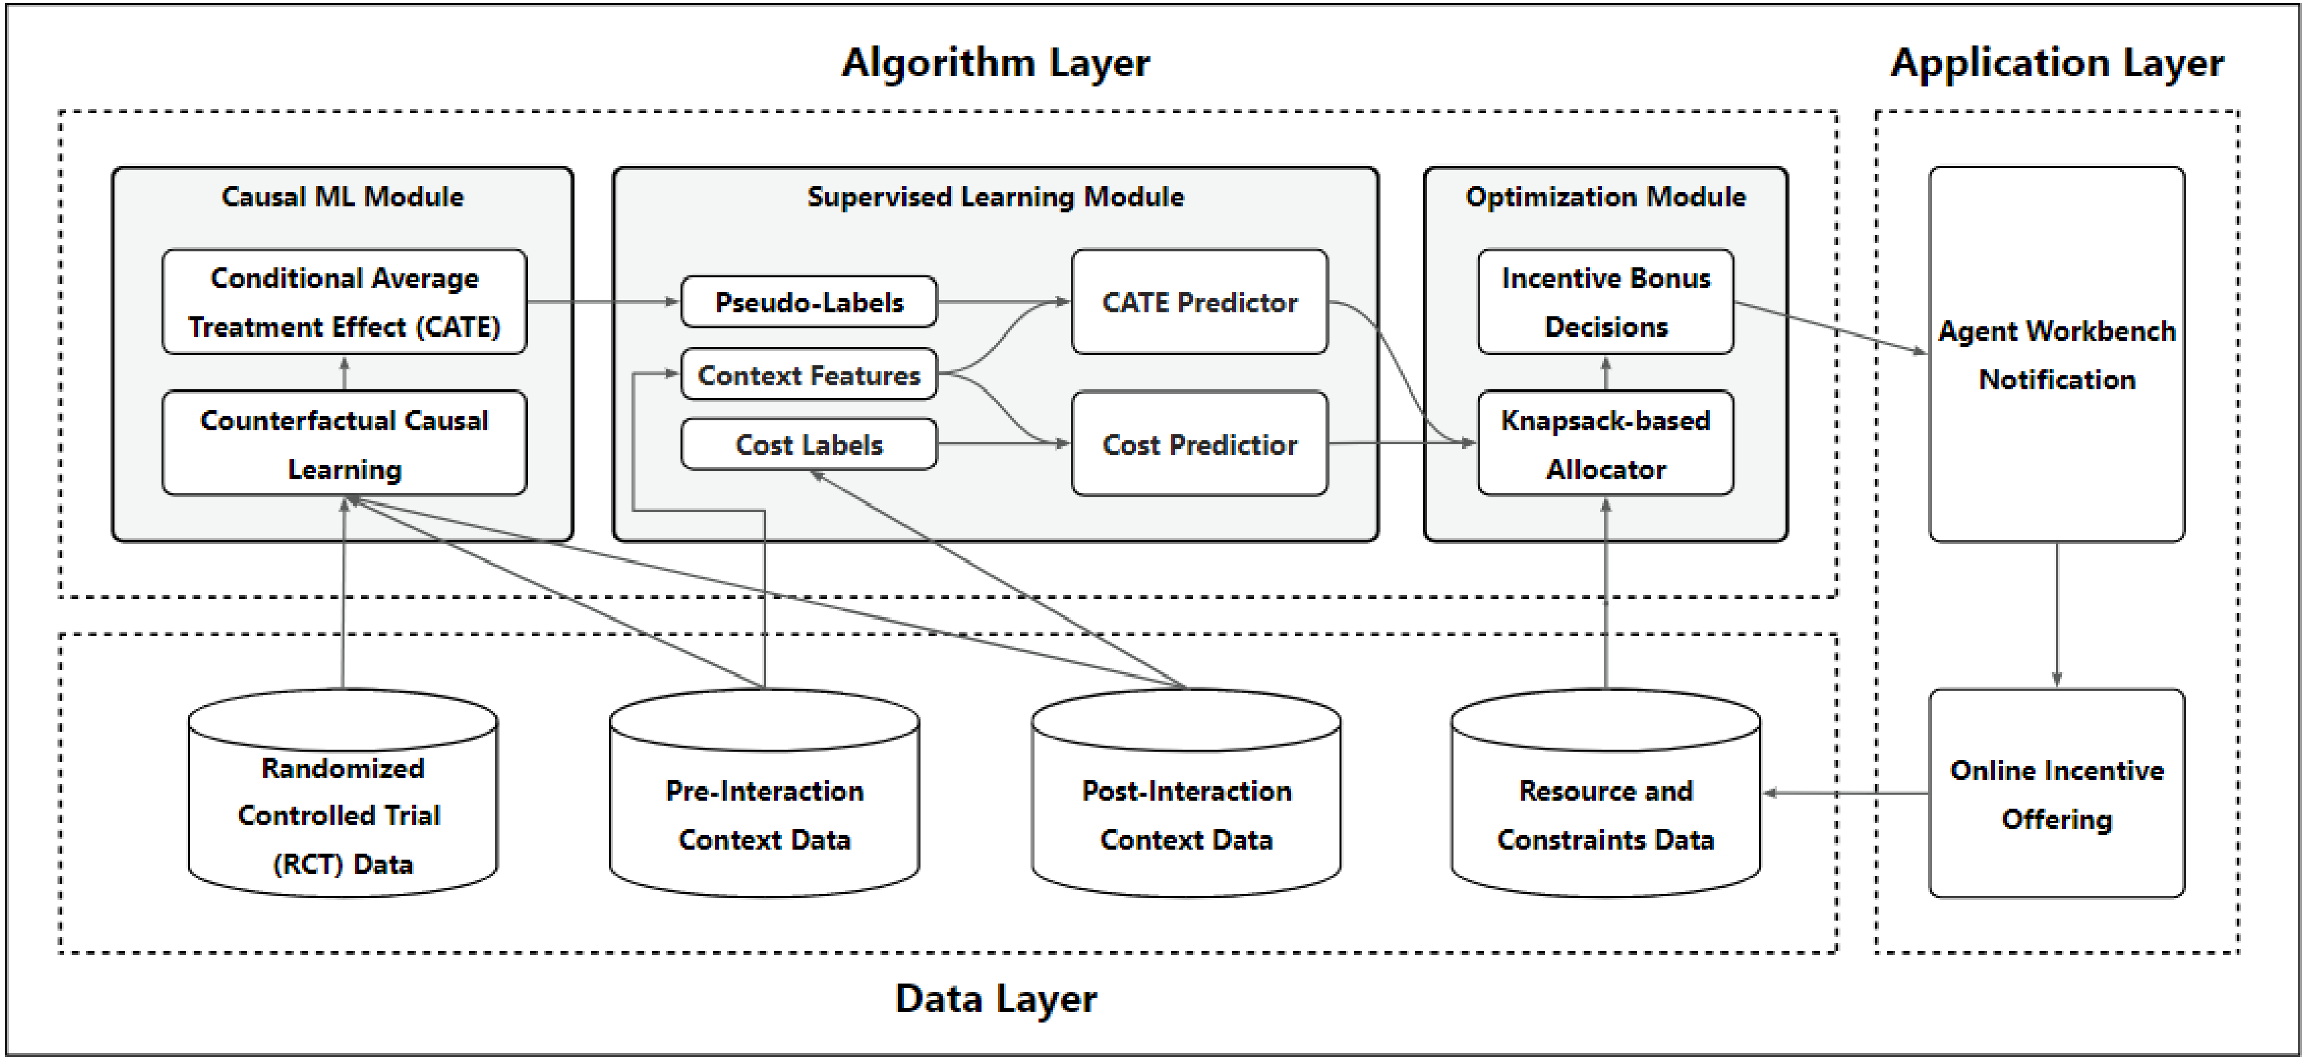

Foundation: RCT Data from a Field Experiment

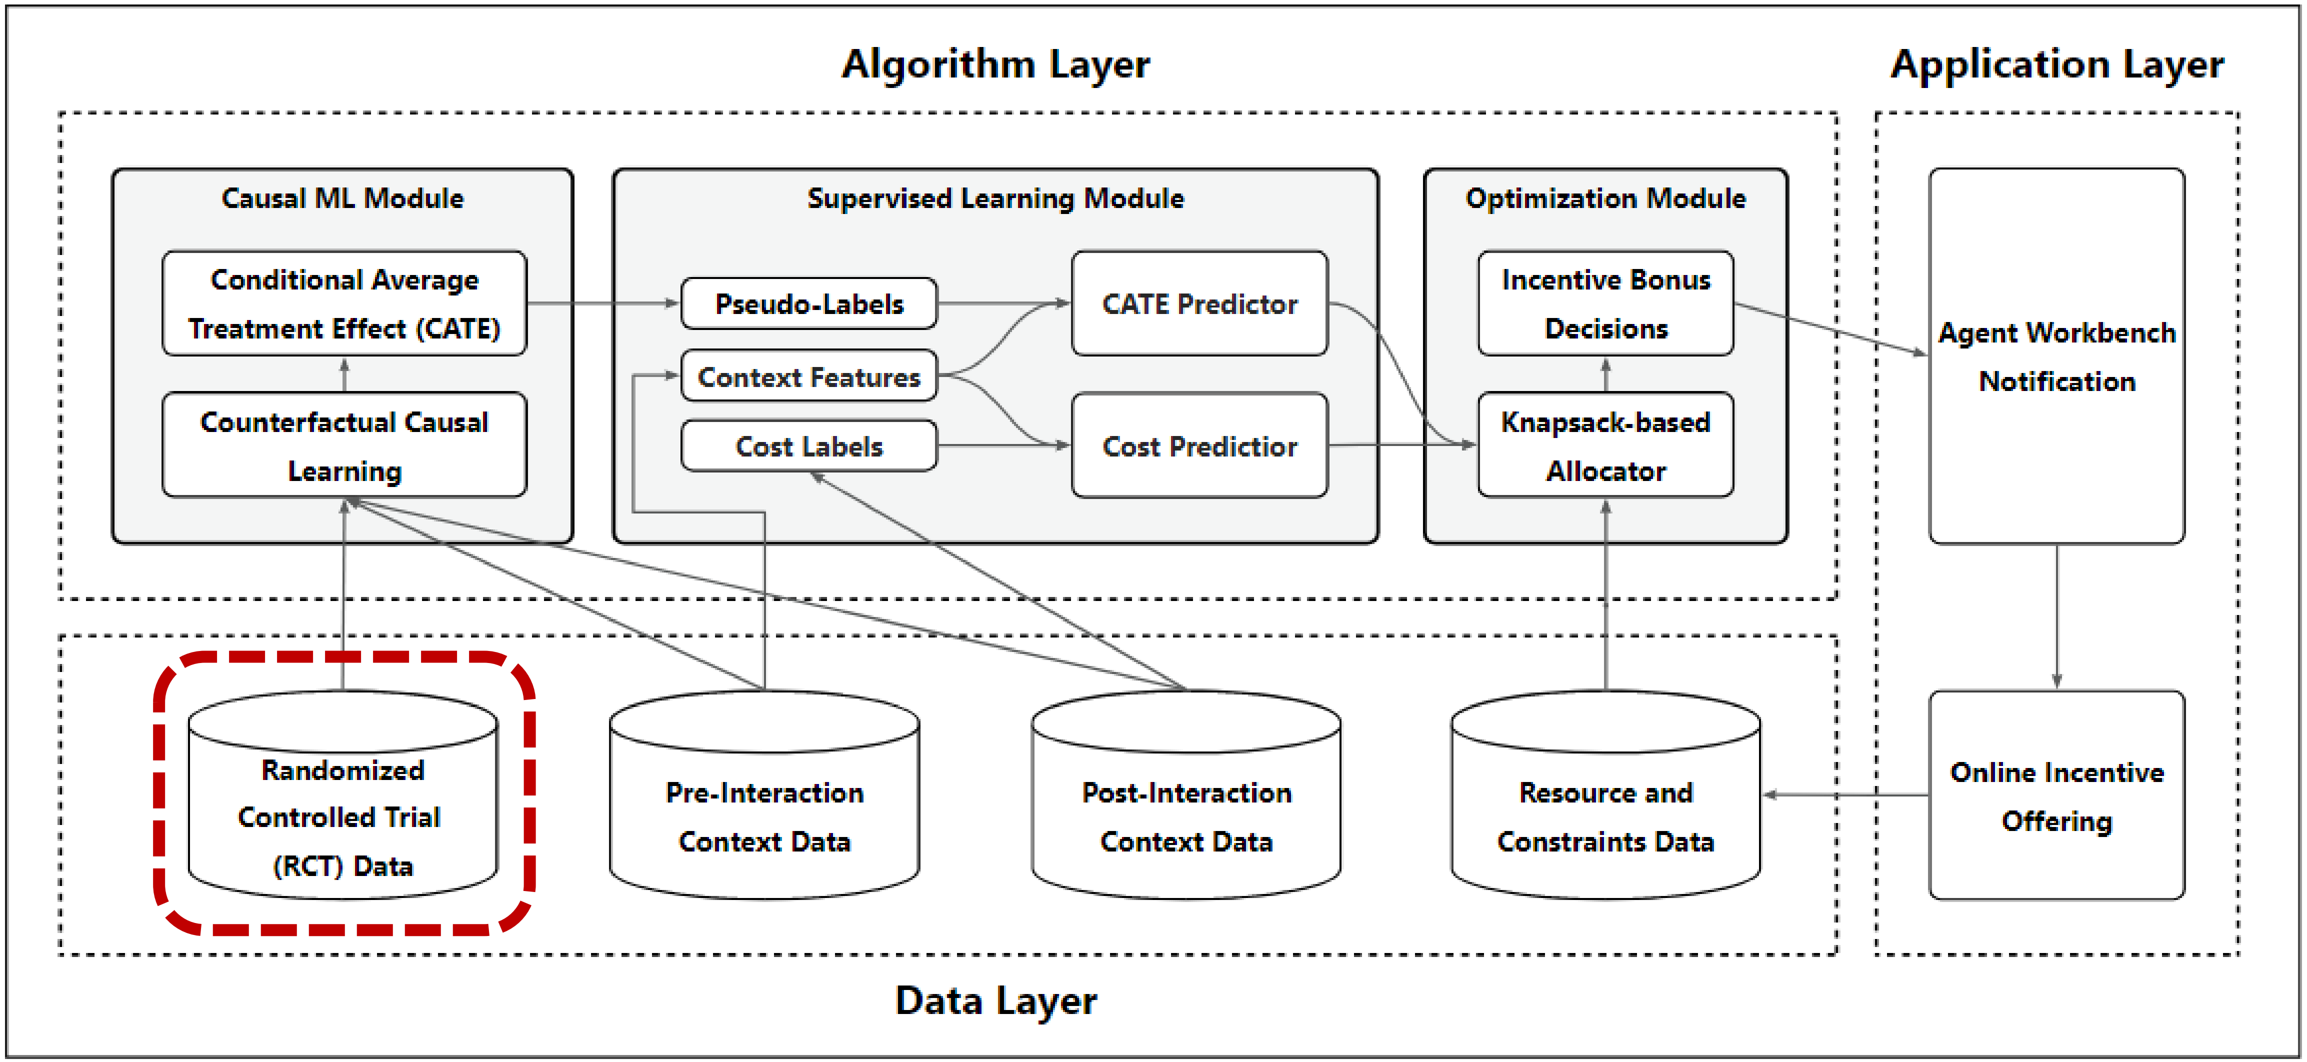

Causal ML Module: Contextualizing Bonus Effects

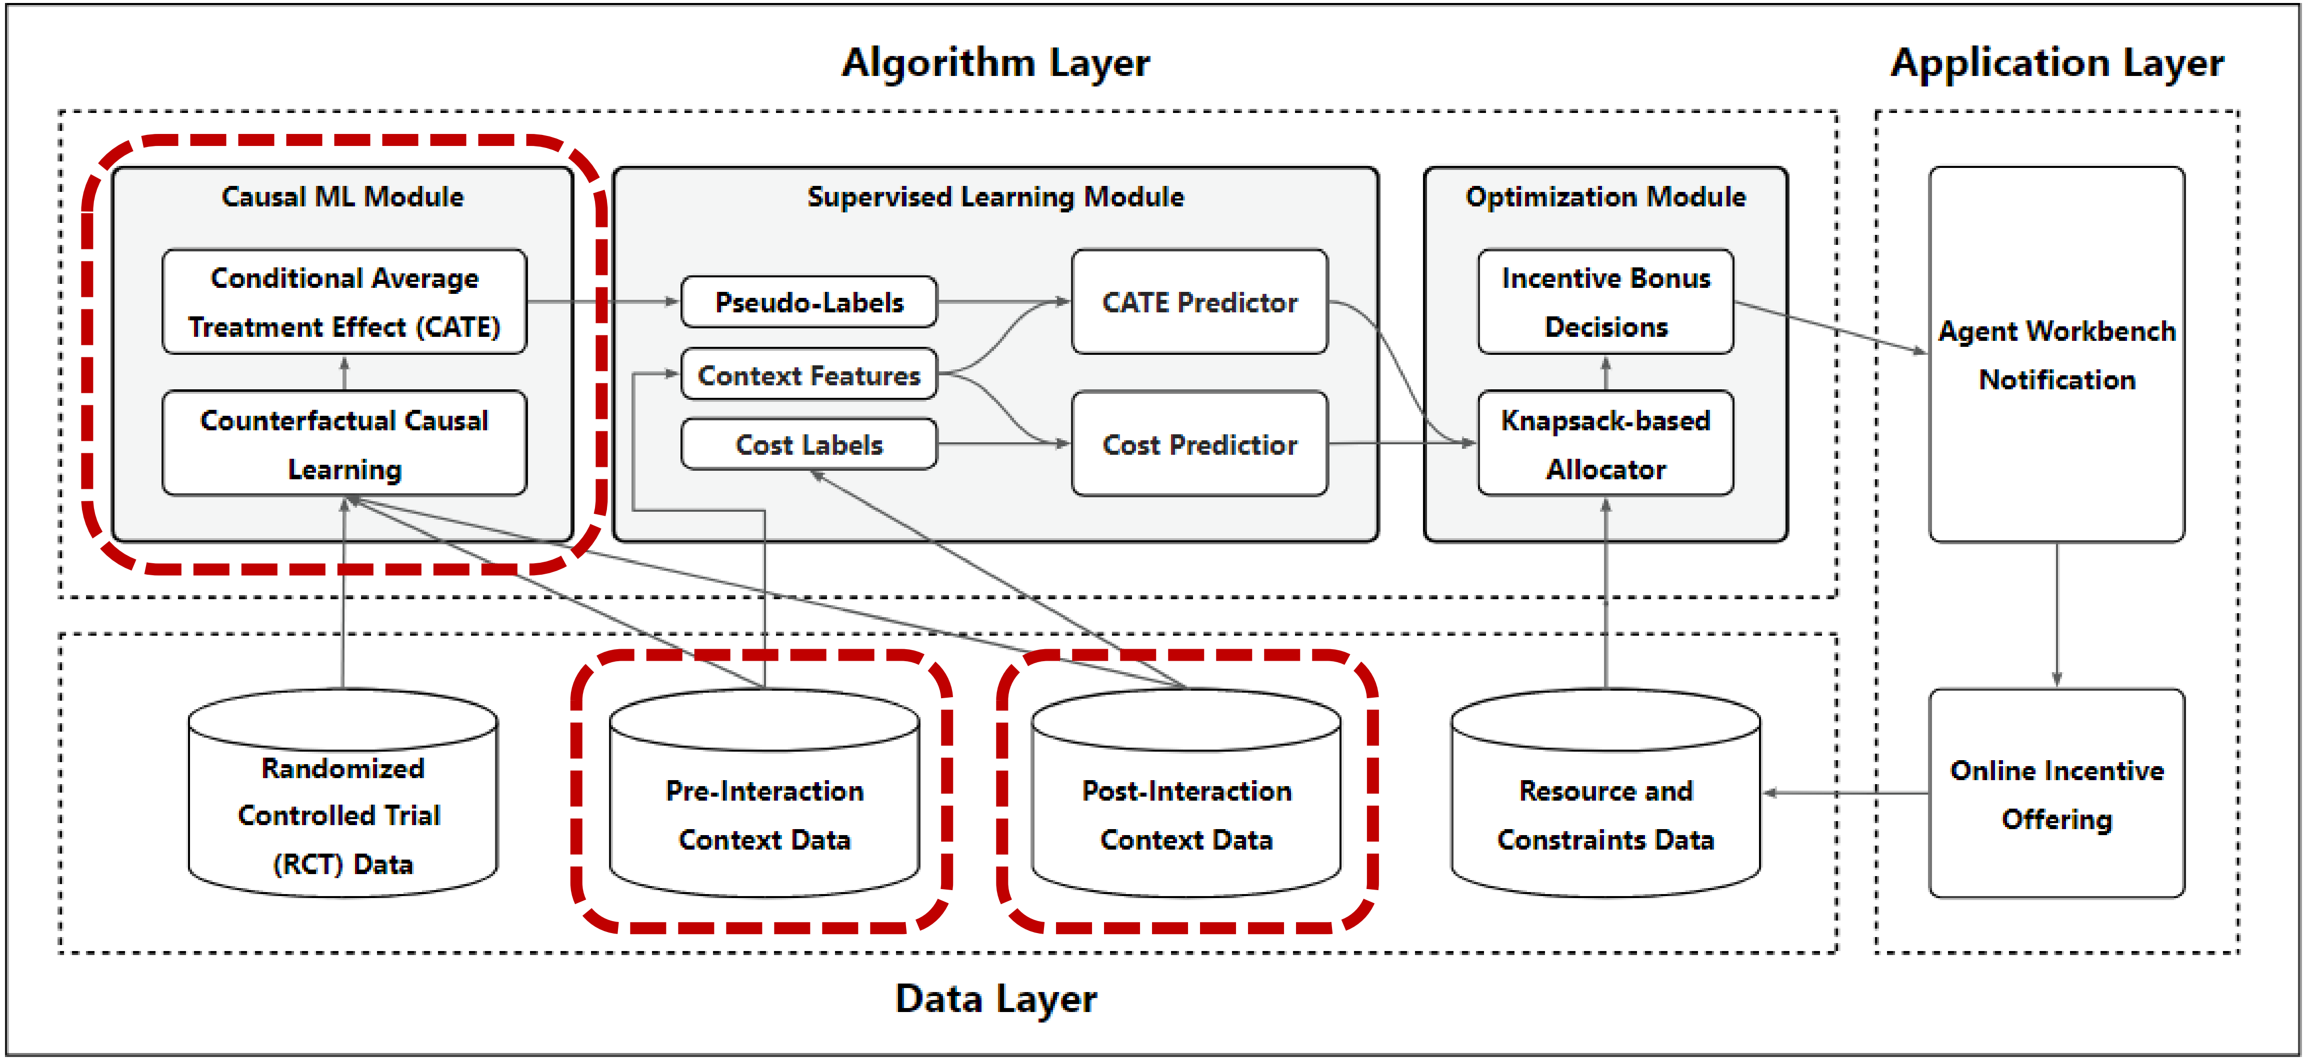

Causal ML Module: Contextualizing Bonus Effects

- **Key Idea**: Use flexible ML predictors to represent the data generation process of an underlying structural causal model
- **Problem Formulation:**
  - Case: $i \in I$
  - Contextual Features: $X_i = \left(X_i^{\text{pre}}, X_i^{\text{post}}\right)$ (pre-chat and post-chat features)
  - Satisfaction Outcome: $Y_i \in \{0, 1\}$ (whether the customer is satisfied)
  - Treatment: $T_i \in \{0, 1\}$
  - Conditional Bonus Effects (CATE):

$$\tau_i(x_i) = \mathbb{E}\left[Y_i \mid X_i = x_i,\, T_i = 1\right] - \mathbb{E}\left[Y_i \mid X_i = x_i,\, T_i = 0\right]$$

- **Key Benefits**: Enables the estimation of case-level bonus effects conditional on high-dimensional *context features*

Supervised ML Module:
Deployable Counterfactual Predictions

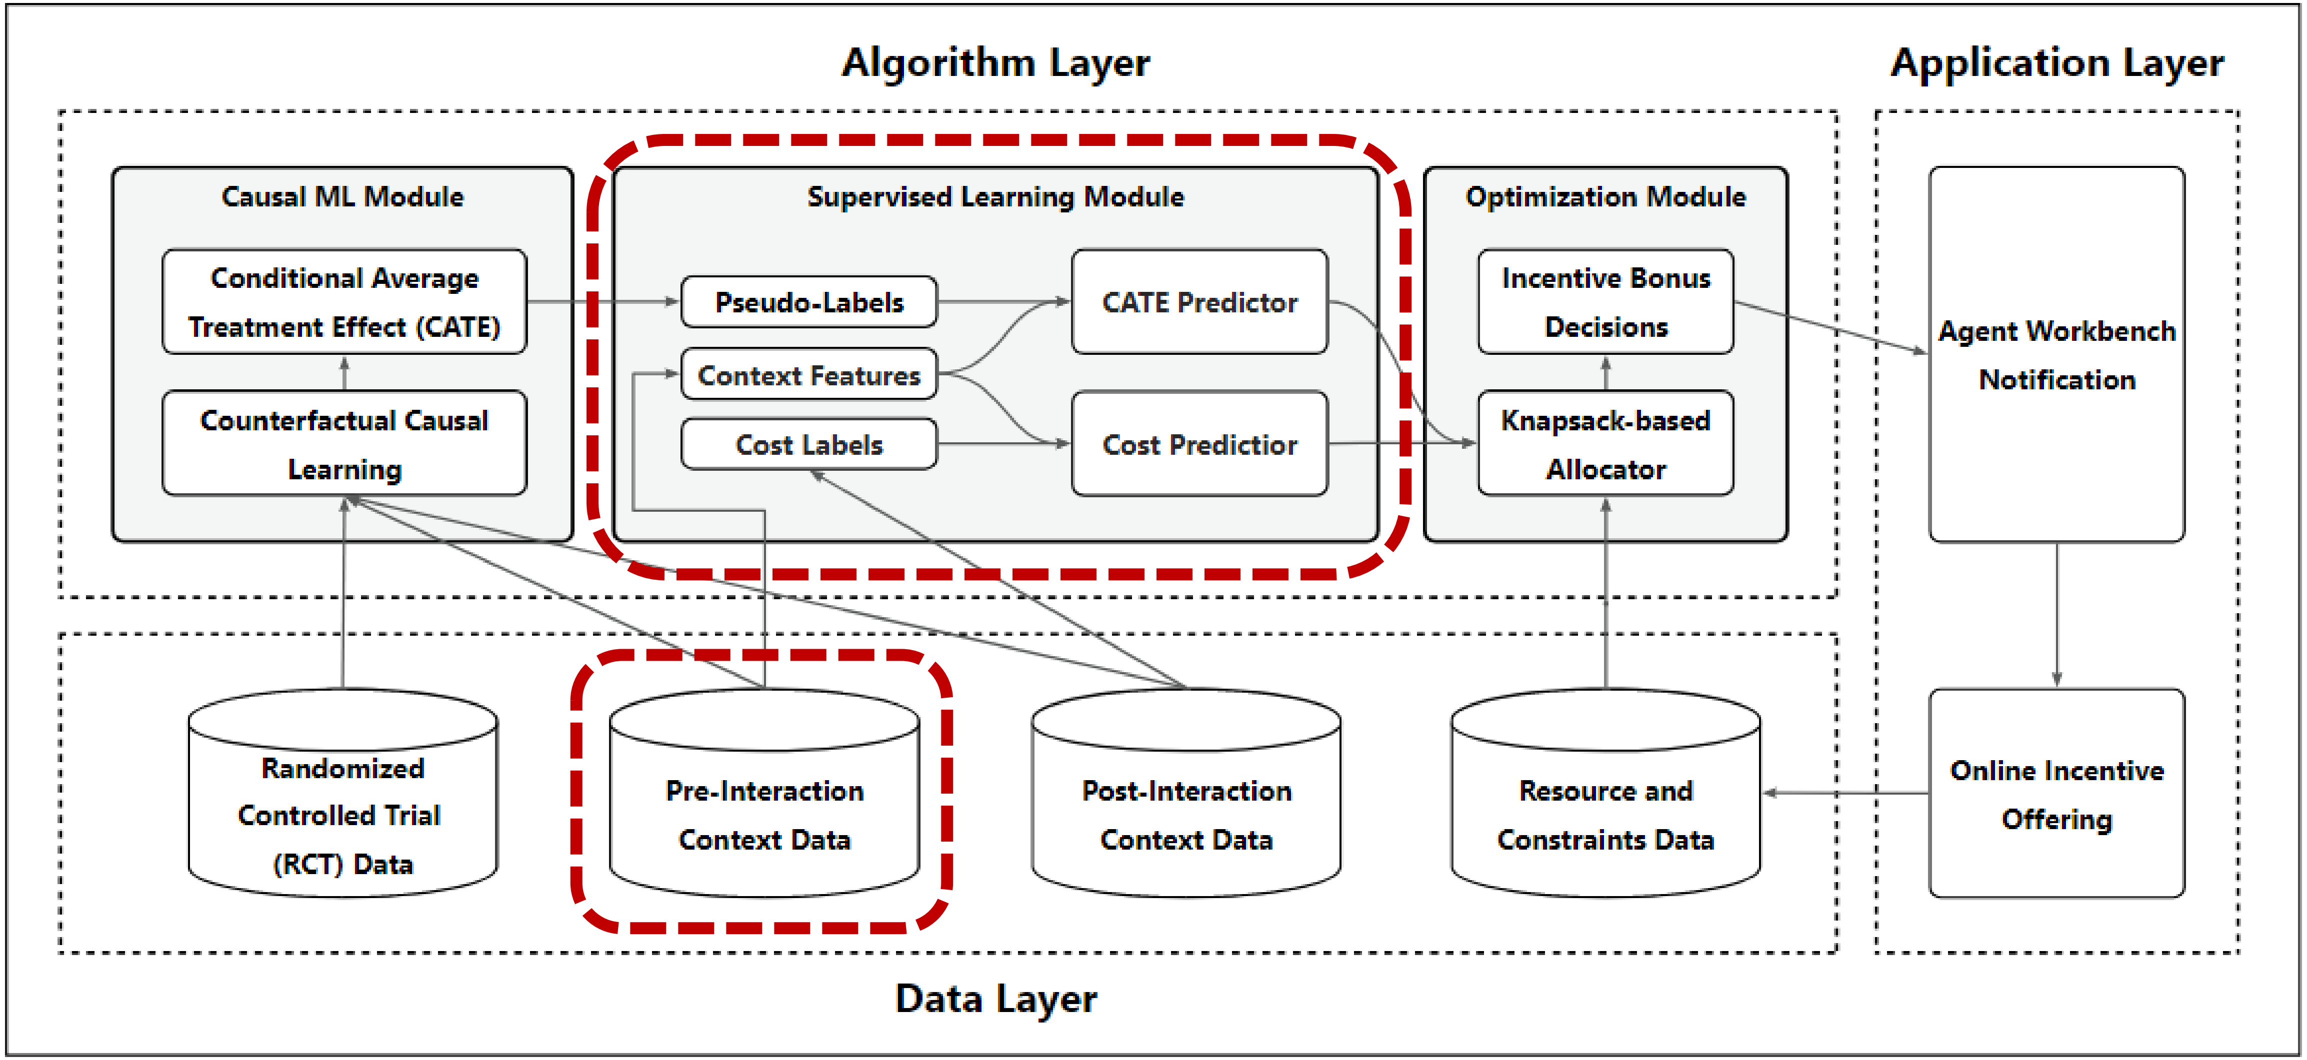

Supervised ML Module: Deployable Counterfactual Predictions

- **CATE Predictor:** $\tau_i = f(X_i^{\text{pre}})$
  - Label: Case-level bonus effects (i.e., output from Causal ML models)
  - Predictor Variables: Pre-interaction context features

- **Cost Predictor:** $c_i = g(X_i^{\text{pre}})$
  - Label: Case-level incentive costs
  - Predictor Variables: Pre-interaction context features

- **Key Benefits**
  - Supports any type of supervised ML models
  - Enables predictions of bonus effects for any given case *before* the chat

Optimization Module:
Acting on Predicted Effects of Bonuses

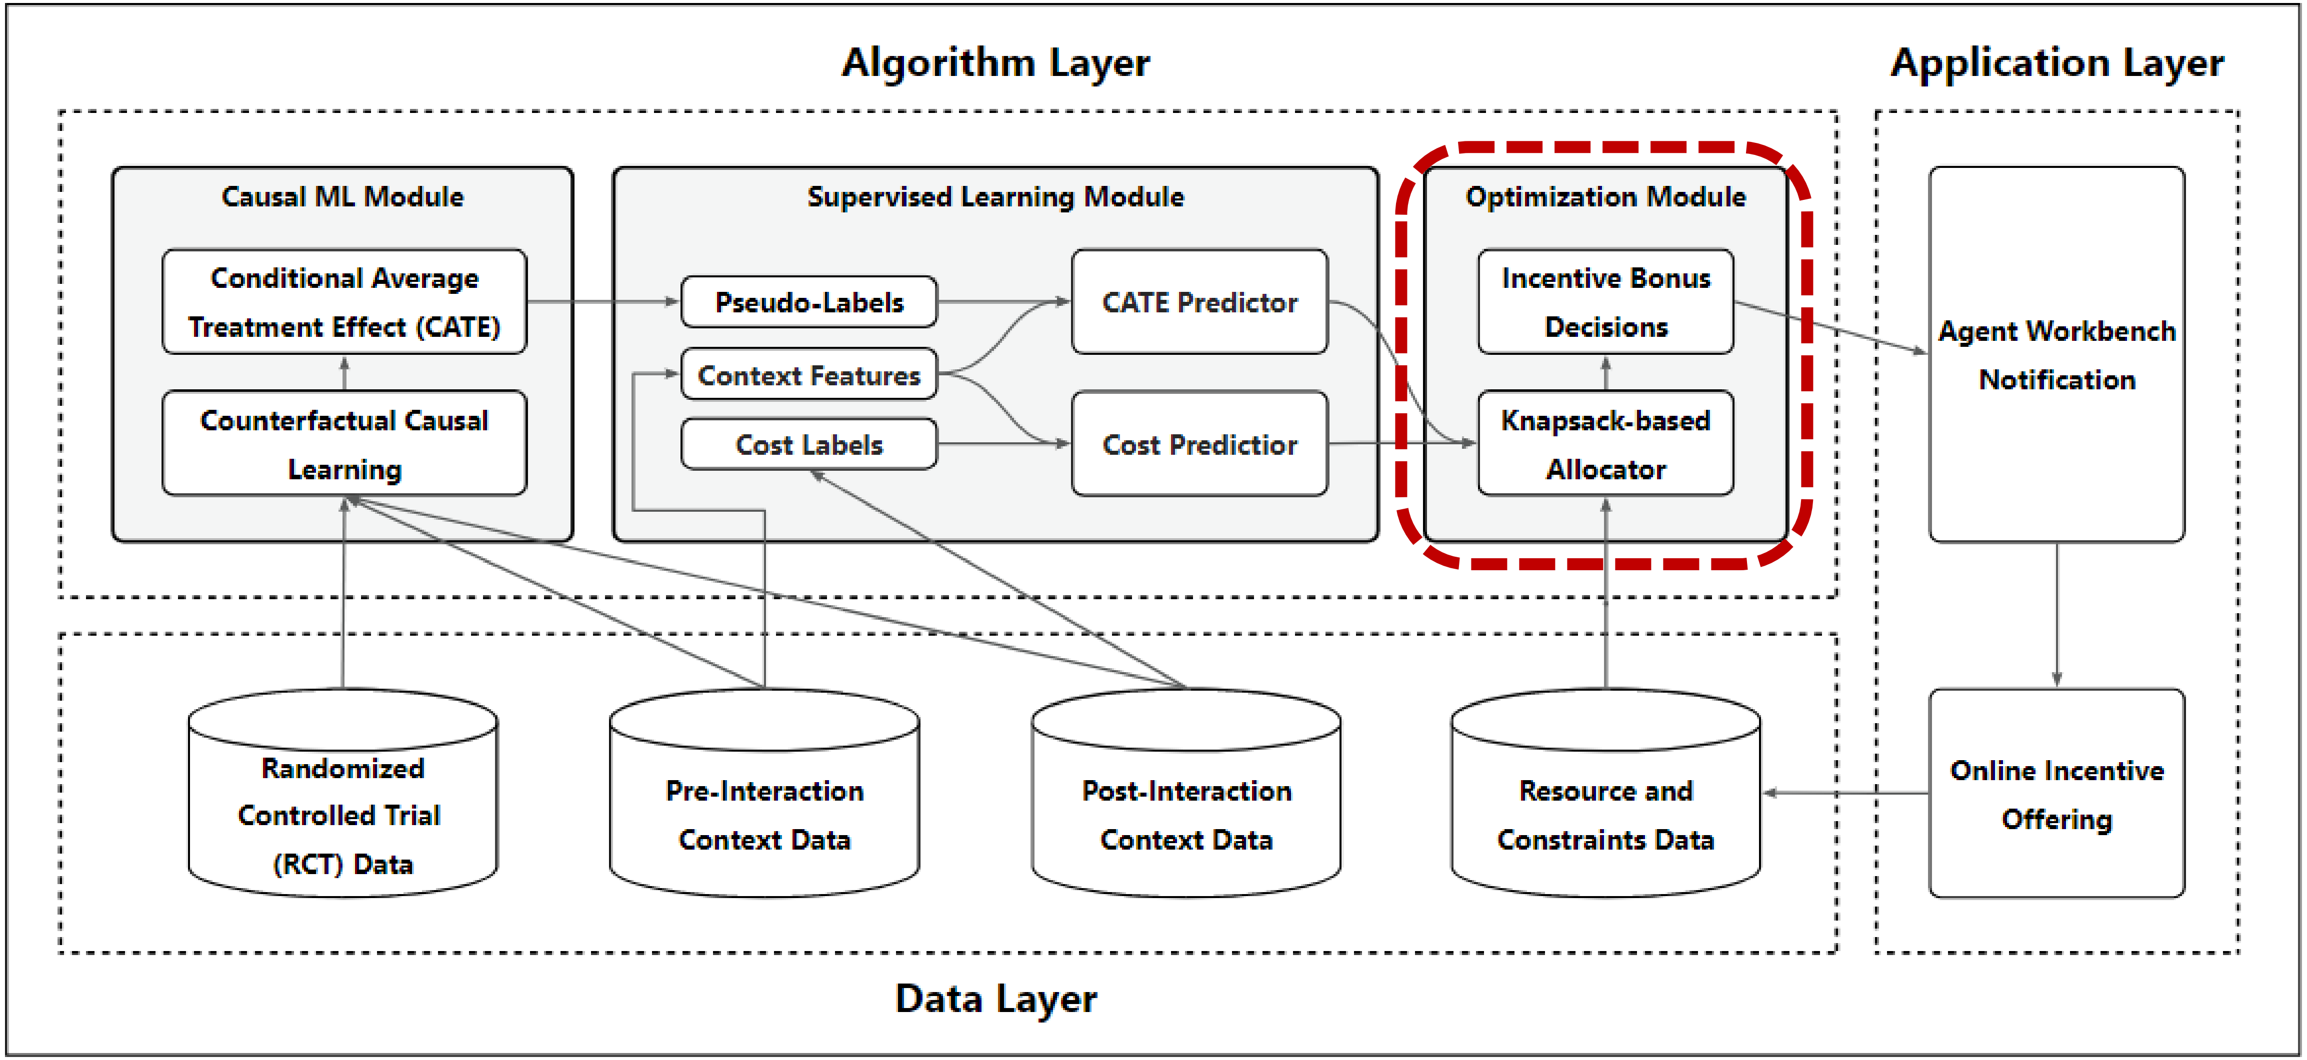

Optimization Module:
Acting on Predicted Effects of Bonuses

Problem Formulation

- Objective: Maximizing the total satisfaction uplift
- Case: $i \in I$
- Bonus Decision: $w_i \in \{0, 1\}$
- Bonus Cost: $\widehat{c}_i$
- Budget at the time of decision: b_0
- Effect of Bonus: $\widehat{\tau}_i \in [\bar{\tau}_i - \delta_i, \bar{\tau}_i + \delta_i]$
- Protection Function: $\Delta(w, \Gamma)$

Uncertainty in uplift predictions

Worst-case impact of uplift uncertainty

$$
\begin{aligned}
z^R(\Gamma) := \quad & \max_w \sum_{i \in I} \bar{\tau}_i w_i - \Delta(w, \Gamma) \\
\text{s.t.} \quad & \sum_{i \in I} \widehat{c}_i w_i \leq b_0 \\
& w_i \in \{0, 1\} \quad \forall i \in I \\
\text{where} \quad & \Delta(w, \Gamma) := \max_s \sum_{i \in I} \delta_i w_i s_i \\
\text{s.t.} \quad & \sum_{i \in I} |s_i| \leq \Gamma \\
& |s_i| \leq 1 \quad \forall i \in I
\end{aligned}
$$

Robustness level

How much uplift deviation is considered for each case

Policy Evaluation

Challenge: The counterfactual outcomes of satisfaction is not observable

Solution: We separate a holdout sample from the randomized field experiment to compare the performance of our approach against several benchmark targeting policies

- Given a policy $\pi : \mathcal{X} \to \{0, 1\}$ and the estimated propensity score $\hat{e}_i$
- Calculate the inverse propensity score-weighted (IPS) estimator of mean satisfaction:

$$\hat{R}_{\mathrm{IPS}}(\pi) = \frac{1}{N} \sum_{i \in I} \left[\frac{1 - w_i}{1 - \hat{e}_i} \left(1 - \pi(x_i)\right) y_{i,T=0} + \frac{w_i}{\hat{e}_i} \pi(x_i)\, y_{i,T=1} \right]$$

Benefits:

- Cost advantages: Enables evaluation of an arbitrary number of targeting policies using only one randomized sample
- Unbiased estimation: The IPS estimator provides an unbiased estimate of the expected satisfaction if the proposed policy had been implemented.

Results

Expected Satisfaction Rate (ESR)

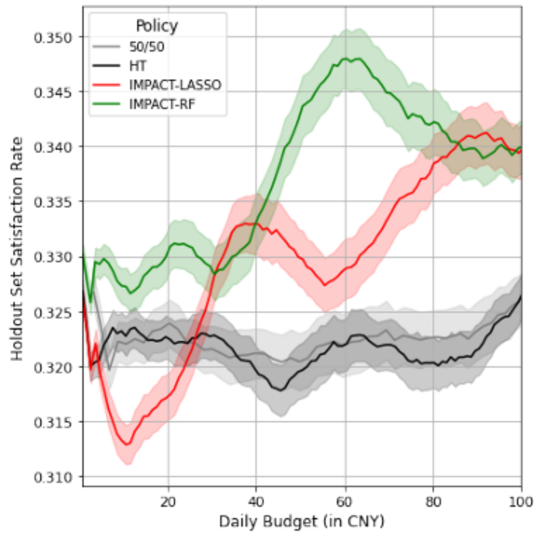

(a) ESR by Daily Budget

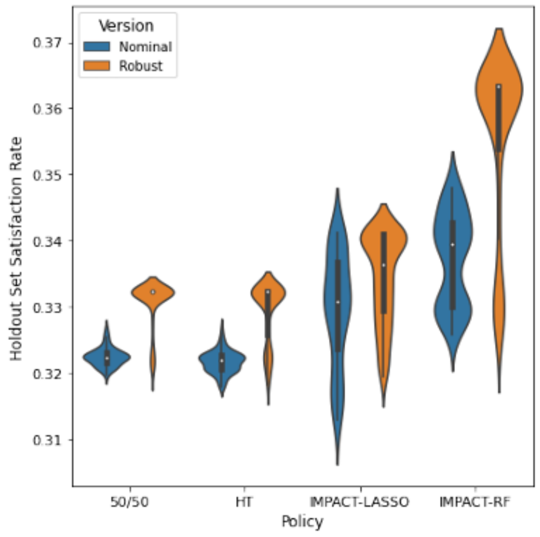

(b) ESR by Model Robustness

Area Under the Cost Curve (AUUC)

Policy	AUCC	Improvement
HT	0.2114 ± 0.0009	/
50/50	0.1770 ± 0.0254	-16.27%
IMPACT-LASSO	0.5228 ± 0.0276	147.30%
IMPACT-RF	0.5455 ± 0.0287	158.04%

Key Insights

- IMPACT consistently outperforms context-free benchmarks, without increasing total incentive spending
- Policies derived from the robust model generally achieve higher satisfaction rates than the nominal model

Conclusions

- We develop a model-based framework (i.e., IMPACT) to deliver cost-effective incentives in service operations.
- Our system is “smart” because it
  - enables context-aware, individualized, proactive bonus targeting
  - adapts to large-scale, high-frequency decision-making settings
  - demonstrates significant value gain compared to the rule-based benchmark