UDC 659.113.71 DOI: https://doi.org/10.64076/iedc251005.04

From playbooks to code: translating bidding strategy into rigorous objective functions and constraints

Igor Ivitskiy
PhD, Associate Professor, Founder,
Doctor Ads LTD, London

Abstract. *The paper formalizes how advertising playbooks translate into mathematical optimization programs for automated bidding. It demonstrates that targets such as budget caps, ROAS floors, and CPA ceilings can be expressed as objective functions and constraints suitable for online primal–dual algorithms. The proposed framework connects common bidding strategies with convex optimization, ensuring stability, fairness, and near-optimal regret under real-time auction dynamics. Special attention is given to the economic interpretation of dual variables as shadow prices that guide budget pacing and ROAS control, bridging heuristic playbooks with rigorous algorithmic implementation.*
Keywords: *automated bidding, ROAS optimization, constrained learning, budget pacing.*

Advertising playbooks often specify budget caps, a desired ROAS, and a tolerance for CPA drift. These phrases only become operational when translated into decision variables, objective functions, and constraints that a bidder can optimize every auction. Let the decision at impression i be a bid b_i, which induces a win probability $w_i(b_i)$, a cost random variable C_i if won, a conversion indicator Y_i, and a conversion value V_i. The campaign budget is B. A simple value-maximizing formulation is:

$$\max_{\{b_i\}} \sum_{i=1}^{N} E[V_i w_i(b_i)] \quad s.t. \quad \sum_{i=1}^{N} E[C_i w_i(b_i)] \le B. \tag{1}$$

Here $E[\cdot]$ denotes expectation over auction and user uncertainty; V_i is the monetary value assigned to the conversion if it occurs; and C_i is the clearing cost conditional on winning. This maps directly to "maximize conversion value within budget", which mirrors platform documentation for Maximize Conversion Value and Target ROAS strategies [1]. To enforce a target ROAS r^*, replace the ratio with a linear constraint using an epigraph trick: the policy must satisfy:

$$\sum_{i=1}^{N} E[V_i w_i(b_i)] - r^* \sum_{i=1}^{N} E[C_i w_i(b_i)] \ge 0. \tag{2}$$

This transforms "keep ROAS at or above r^*" into a form that primal-dual methods can handle with online mirror descent updates for the dual variables, which act as shadow prices on spend and on the ROAS gap [2]. If a playbook instead specifies a CPA ceiling τ, constrain the program by

$$\sum_{i=1}^{N} E[C_i w_i(b_i)] - \tau \sum_{i=1}^{N} E[Y_i w_i(b_i)] \le 0, \tag{3}$$

where Y_i equals 1 when a conversion occurs and 0 otherwise. This expresses "cost per action at or below τ" without resorting to unstable post-hoc target tweaks [4].

Economic operations require time consistency. Let $\mathcal{T}_t$ denote auctions in time bucket t and B_t the pacing budget. Smooth delivery becomes a set of intertemporal constraints

$$\sum_{i \in \mathcal{T}_t} E[C_i w_i(b_i)] \leq B_t, \quad t = 1, \dots, T. \tag{4}$$

Lagrange multipliers on these constraints yield time-varying bid multipliers that track budget digestion. In practice this reproduces the well-known “pacing parameter” that declines when under-spending and rises when overspending, but here it is grounded in a convex program with measurable dual gaps [2]. The Lagrangian for value maximization with a global budget and a ROAS floor introduces dual variables $\lambda \geq 0$ for spend and $\mu \geq 0$ for the ROAS constraint:

$$\begin{aligned} L(b, \lambda, \mu) &= \sum_i E[V_i w_i(b_i)] - \lambda \left(\sum_i E[C_i w_i(b_i)] - B \right) \\ &- \mu \left(r^* \sum_i E[C_i w_i(b_i)] - \sum_i E[V_i w_i(b_i)] \right). \end{aligned} \tag{5}$$

Optimizing bids b_i while updating λ and μ online yields a principled rule: increase the effective bid when the marginal value exceeds the dual-weighted marginal cost, and reduce otherwise; update the duals to keep spend and ROAS on target [2]. This unifies the folk rule “lower tROAS to grow volume” with a formal statement: lowering r^* relaxes the ROAS constraint, which reduces μ in steady state and expands the feasible set of auctions the bidder will enter.

A second translation layer addresses what the value term represents. Value-based bidding treats V_i as the advertiser’s measured revenue or contribution margin, which aligns with maximize-value programs under ROAS constraints. When the strategic goal is to maximize incremental outcomes rather than attributed outcomes, define V_i as the lift $E[\,Yi \mid ad\,] - E[\,Yi \mid no\ ad\,]$. The corresponding “lift-based” objective bids on causal effect, not propensity, which can increase true actions even at the same spend [5]. Recent engineering work shows that unbiased lift predictors, combined with a pacing multiplier α\alphaα and standard CTR models, can be deployed under first-price auctions and CPC billing while preserving DSP profitability [5]. At portfolio level, offline constrained reinforcement learning can learn budget allocation policies from logged data that satisfy long-term constraints and converge, which gives an implementable path to cross-channel spend programs [3].

Playbooks also demand stability and fairness. To reduce oscillations while targets change, add a stability penalty $\Omega(\Delta b)$ on bid changes relative to a rolling baseline $\overline{b_i}$:

$$\Omega(\Delta b) = \gamma \sum_{i=1}^{N} \left(b_i - \overline{b_i}\right)^2, \tag{6}$$

where $\gamma \geq 0$ controls aggressiveness.

If the business requires minimum exposure for segment g, define $\mathcal{G}_g$ and enforce share constraints with indicator functions $1\{\cdot\}$:

$$\sum_{i=1}^{N} 1\{i \in \mathcal{G}_g\} E[C_i w_i(b_i)] \geq s_g B, \quad (7)$$

where s_g is the spend share floor for the segment and $1\{\cdot\}$ equals 1 if impression i belongs to segment g and 0 otherwise. Risk constraints enter as bounds on variance of daily ROAS or as percentile pacing limits, both of which admit dualized treatments akin to budget pacing [6].

Two implementation notes bridge math and platform controls. First, portfolio strategies that share targets across campaigns are exactly the multi-budget, multi-constraint case. Use one set of duals per constraint and a global objective on portfolio value. Second, the "knobs" exposed by platforms (target ROAS, target CPA, budget caps, device exclusions) can be treated as parameters that set constraints or bounds in the program. Raising the ROAS target tightens the ratio constraint, which increases the effective shadow price μ; capping device traffic introduces additional sub-budget constraints by inventory slice; and portfolio bid strategies operationalize shared duals across campaigns [3]. Clear mappings help explain why multi-predictor bidders that combine CTR with win-price models often outperform pure propensity bidders under the same budget, because the objective maximizes expected clicks or value subject to binding cost constraints that depend on the winning price distribution [6].

Finally, the literature supports the convergence and feasibility of these programs under realistic assumptions. Online primal-dual algorithms can respect return-on-spend and budget constraints while achieving near-optimal regret, which formalizes "stay on target without starving volume" [2]. Constrained CPA programs formulated over learned bid-landscapes also deliver policy-compliant recommendations in practice [4]. Lift-based variants can raise true incremental actions when attribution is noisy, which is economically relevant in privacy-constrained environments [5]. In sum, turning playbooks into code means declaring the objective and constraints, linearizing ratios, and letting dual variables act as interpretable control signals. The outcome is a bidder that optimizes what the business actually wants, not what the interface merely reports [1-6].

References

1. Feng Z., Padmanabhan S., Wang D. Online Bidding Algorithms for Return-on-Spend Constrained Advertisers. In: *Proceedings of the ACM Web Conference 2023*. New York: ACM, 2023. DOI: 10.1145/3543507.3583491.

2. Cai T., Jiang J., Zhang W., Zhou S., Song X., Yu L., Gu L., Zeng X., Gu J., Zhang G. Marketing Budget Allocation with Offline Constrained Deep Reinforcement Learning. In: *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining (WSDM '23)*. New York: ACM, 2023. DOI: 10.1145/3539597.3570486.

3. Ivitskiy, I. Mastering Google Ads Portfolio Bid Strategies. *The Doctor Ads* (blog). Available at: https://blog.thedoctorads.com/mastering-google-ads-portfolio-bid-strategies/ (accessed 18 Sep 2025).

4. Kong D.; Shmakov K.; Yang J. Demystifying Advertising Campaign Bid Recommendation: A Constraint target CPA Goal Optimization. *arXiv preprint* arXiv: 2212.13915, 2022. DOI: 10.48550/arXiv.2212.13915.

5. Moriwaki D., Hayakawa Y., Matsui A., Saito Y., Munemasa I., Shibata M. A Real-World Implementation of Unbiased Lift-based Bidding System. *arXiv preprint* arXiv:2202.13868, 2022. DOI: 10.48550/arXiv.2202.13868.

6. Lin C.-C., Chuang K.-T., Wu W. C.-H., Chen, M.-S. Budget-Constrained Real-Time Bidding Optimization: Multiple Predictors Make It Better. *ACM Transactions on Knowledge Discovery from Data*, 2020, 14(2): 1–27. DOI: 10.1145/3375393.

UDC 338:502:342:351

Systematization of the integrated assessment of the ecological and economic mechanism of protected areas in the context of environmental policy implementation

Maksym Horinshtein
Institute of Agroecology and Environmental Management
of the National Academy of Agrarian Sciences of Ukraine, Kyiv
Mariya Vysochanska
Institute of Agroecology and Environmental Management
of the National Academy of Agrarian Sciences of Ukraine, Kyiv
Wojciech Walat
University of Rzeszów, Rzeszów

Abstract. *The article explores the systematization of the integrated assessment of the ecological and economic mechanism of protected areas within the framework of environmental policy implementation. The study emphasizes the importance of combining economic and environmental instruments–such as taxes, fees, fines, and monitoring systems–to ensure the sustainable use of natural resources. The proposed systematization provides a comprehensive structure for evaluating the interaction between environmental and economic processes, including the development of indicators, the use of integral indices, and the application of environmental assessment tools. The results underline that effective environmental management requires coordination of state policy, economic incentives, and social responsibility to achieve sustainable development goals.*
Keywords: *ecological and economic mechanism, systematization, sustainable development, environmental policy, protected areas.*