

Optimal Mean-Field Stackelberg Game-Inspired Resource Allocation for Scalable Multi-Modal Semantic Communications in 6G Networks

Tung Son Do, Thanh Phung Truong, Manh Cuong Ho, Quang Tuan Do, Nhu-Ngoc Dao, Hyunbum Kim, and Sungrae Cho

Abstract—Sixth-generation networks are expected to support large-scale multimodal semantic communication services under stringent spectrum constraints. Most current frameworks ignore semantic-level quality heterogeneity across modalities or depend on centralized optimization, which becomes computationally prohibitive when serving massive user populations. To address these limitations, this work proposes a mean-field Stackelberg equilibrium (MFSE) framework that applies the mean-field game theory and Stackelberg pricing to optimize scalable resource allocation and modality-aware semantic quality jointly. Acting as the Stackelberg leader, the base station sets modality-specific prices to maximize the network revenue and semantic welfare. Each user independently optimizes its per-modality bandwidth allocation against the mean-field aggregate demand, subject to minimum quality-of-service requirements. The leader pricing is optimized using a cosine-annealed Adam optimizer with Polyak price averaging and exact analytical gradients derived via the implicit function theorem. This work proves the existence and uniqueness of the MFSE, and the resulting strategy profile forms an approximate Nash equilibrium for the finite-player game. Simulations show that the proposed framework achieves 98.89% of the centralized semantic welfare and, under task-aware steering, attains the highest literal end-to-end task score. Furthermore, the equilibrium computation is $O(1)$ in the user population, while distributed deployment scales as $O(NJ)$.

Index Terms—Mean-field game (MFG), multimodal, resource allocation, semantic communication (SemCom), Stackelberg game.

I. INTRODUCTION

THE accelerated increase in intelligent services (e.g., immersive extended reality and autonomous systems and the Internet of Senses) is challenging next-generation wireless networks and demands new communication paradigms beyond Shannon’s classical bit-level transmission. To address this need, semantic communication (SemCom) extracts and transmits only task-relevant meaning [1], reducing bandwidth while preserving task performance. In multimodal SemCom, weighted semantic quality (WSQ) measures the weighted

task utility aggregated across modalities, users, and resource-allocation decisions. Real-world applications (telehealth, autonomous driving, and immersive conferencing) simultaneously involve multiple data modalities; therefore, the field has progressed from unimodal systems to multimodal architectures, jointly addressing text, images, and speech [2]. One challenge is that each modality has distinct semantic quality characteristics and heterogeneous resource requirements, complicating unified resource allocation.

Multimodal SemCom resource allocation faces three challenges distinct from conventional radio resource management. First, modality heterogeneity indicates that text, images, and speech display different resource-quality relationships. In this work, these relationships are not treated as unvalidated proxies; they are calibrated to measured task metrics from DeepSC text transmission, Deep joint source–channel coding image transmission, and DeepSC-S speech transmission. Therefore, a single allocation policy cannot effectively serve all modalities, and per-modality strategies must be jointly optimized under a fixed total bandwidth budget. Second, user heterogeneity arises because applications assign different semantic priorities to each modality (e.g., telehealth prioritizes image fidelity, whereas conferencing favors speech clarity), creating a combinatorial challenge in any finite-player game formulation. Third, semantic coupling arises from shared neural encoders, content-importance weights, cross-modal correlation, and congestion at the base station (BS). Although codec architecture design is outside the scope of this paper, this coupling is a first-class resource-allocation feature rather than an implementation detail. Consequently, the problem is a strategic game rather than an independent optimization process.

The existing approaches fail to address these challenges jointly. Centralized optimization methods and training-heavy deep reinforcement learning (DRL) schemes [3], [4] can provide strong empirical policies, but they often require repeated data collection, retraining, or centralized problem information as the population and preference distribution change. Moreover, game-theoretic formulations offer equilibrium guarantees; however, existing studies are limited in scope.

For example, Güler *et al.* [5] addressed single-word transmissions without resource allocation, and Liu *et al.* [6] applied a single uniform price solved using DRL. Further, Aggarwal *et al.* [7] restricted the mean-field game (MFG) [8], [9] model to single-modality binary decisions. Therefore, no current framework delivers scalable and analytically tractable

Corresponding author: Sungrae Cho.

Tung Son Do, Thanh Phung Truong, Manh Cuong Ho, Quang Tuan Do, and Sungrae Cho are with the School of Computer Science and Engineering, Chung-Ang University, Seoul 06974, Republic of Korea (e-mail: {tsdo, tptruong, hmcuong, dqtuang}@uclab.re.kr; srcho@cau.ac.kr).

Nhu-Ngoc Dao is with the Department of Computer Science and Engineering, Sejong University, Seoul, Republic of Korea (e-mail: nndao@sejong.ac.kr).

Hyunbum Kim is with the Department of Embedded Systems Engineering, Incheon National University, Incheon, Republic of Korea (e-mail: hyunbumkim@ieec.org).

game-theoretic multimodal resource allocation with continuous per-modality demands, calibrated modality-specific utility, fixed-total bandwidth provisioning, and explicit semantic coupling.

Motivated by these observations, this work proposes a mean-field Stackelberg equilibrium (MFSE) framework that applies the MFG theory and Stackelberg pricing to achieve scalable, distributed, and analytically tractable multimodal SemCom resource allocation with guaranteed equilibrium. The primary contributions of this paper are as follows:

- 1) We have formulated a multimodal SemCom resource allocation problem using the mean-field Stackelberg game theory. The framework is generalized to a scalable scenario with a massive number of users, where the BS sets modality-specific prices, and each user independently optimizes the per-modality bandwidth demands against the mean-field aggregate statistics.
- 2) We have designed modality-specific semantic quality functions that are calibrated to measured task metrics from real semantic codecs as bounded, capacity-compressed, strictly concave surrogates, and introduce a computational congestion model at shared neural encoders. Moreover, this work derives closed-form per-modality bandwidth allocations via Karush–Kuhn–Tucker (KKT) conditions to foster efficient equilibrium computation.
- 3) We have proved the existence and uniqueness of the MFSE via contraction mapping, establish an ε -Nash equilibrium bound of $O(1/\sqrt{N})$, and derive exact gradients via the implicit function theorem (IFT), where the same contraction condition guarantees convergence and well-posedness.
- 4) We have validated the framework via extensive simulations. The results demonstrate that the MFSE achieves nearly centralized and optimal semantic quality, matches or exceeds strong DRL and finite-player baselines across all tested population scales, and maintains $O(1)$ equilibrium computation time independent of the number of users. This $O(1)$ -in- N property holds for the equilibrium (mean-field) computation at a fixed Monte Carlo sample size, whereas deployment and evaluation over all realized users scale as $O(NJ)$. Furthermore, the empirical ε -Nash gap confirms the theoretical $O(1/\sqrt{N})$ convergence rate.

The remainder of this paper is organized as follows. Section II reviews the related work, and Section III presents the system model and problem formulation. Next, Section IV derives the MFSE and proves the approximate Nash property. Then, Section V provides the numerical results. Finally, Section VI concludes the paper.

II. RELATED WORK

A. Game-Theoretic Approaches

Game theory provides a framework for modeling strategic interactions in allocating communication resources. Güler *et al.* [5] formulated SemCom as a Bayesian and sequential game, proving the existence of Bayesian Nash equilibria,

but only for single-word transmissions without allocating resources. Similarly, Xiao *et al.* [10] presented an information-theoretic analysis of strategic SemCom, restricted to point-to-point communication with discrete semantic sources. For allocating multiuser resources, Thomas and Saad [11] introduced the Stackelberg hypergame theory, reducing transmitted bits by about 45%, although without multimodal considerations. In the broader literature, Stackelberg games have been applied to RIS-aided anti-jamming [12], where a hierarchical leader–follower interaction between the legitimate system and the jammer governs the joint power and reflecting-surface control, and to cloud-network convergence resource allocation [13]; however, these frameworks optimize physical-layer anti-jamming or service provisioning rather than modality-aware SemCom quality. Additionally, Moirangthem *et al.* [14] employed Shapley values for fair spectral allocation but addressed only single-modality channel assignment. Fantacci and Picano [15] applied matching theory for allocating single-modality multiuser SemCom resources.

Among studies targeting multimodal SemCom, Yang *et al.* [16] designed a Stackelberg-matching game for secure two-modality SemCom, focusing on physical-layer security rather than general resource optimization. Further, Ji *et al.* [17] formulated a weighted potential game for D2D SemCom but addressed only single-modality text. Notably, Liu *et al.* [6] proposed a Stackelberg game for multimodal SemCom bandwidth pricing in mobile AIGC networks. Although their work is most directly related to that in this paper, their formulation employs a single uniform price without user competition and applies DRL-based solutions that do not scale with the number of users. Moreover, Aggarwal *et al.* [7] applied extended MFGs to multiteam dynamic SemCom with a proven ε -Nash property, yet the work was restricted to single-modality binary decisions.

The MFG theory [8], [9] offers a tractable approach for large-population games by substituting the intractable N -player coupling with a representative agent that interacts with an aggregate population distribution. To solve high-dimensional MFG systems, Zeng *et al.* [18] combined physics-informed operator learning with DRL, showing that neural operators can approximate mean-field solution maps while reducing repeated numerical solves. This line of work motivates learning-aided MFG computation, but it does not address wireless pricing, semantic utility calibration, or distributed bandwidth allocation. The MFG-Stackelberg combination has been applied to allocating wireless resources. For example, Su *et al.* [19] addressed the wireless powered Internet of Things (IoT). Wang *et al.* [20] studied D2D power control, and Benamor *et al.* [21] applied MFG to distributed power control in hybrid nonorthogonal multiple access (NOMA). In particular, Li *et al.* [22] proposed a learning-aided MFG framework for NOMA–mobile edge computing (MEC) resource allocation in ultra-dense networks, sharing the representative-agent scaling principle and learning-aided mean-field updates; nevertheless, it focuses on conventional communication and computation metrics rather than modality-specific semantic utilities, calibrated quality functions, and per-modality congestion coupling. Zhang *et al.* [23] combined MFG with mul-

TABLE I: Comparison with Existing Work

Ref.	Game	Modal.	MU	Scal.	Equil.	Dist.
[5]	Bayes/Seq.	Words	×	—	BNE	—
[6]	Stackelberg	T-I	✓	×	SE	×
[7]	MFG	Single	✓	✓	ϵ -NE	✓
[22]	MFG	NOMA-MEC	✓	✓	MFE	✓
[12]	Stackelberg	PHY	✓	Mod.	SE	Part.
[13]	2-layer Stack.	Cloud	✓	Mod.	SE	Part.
[10]	Stack./Nash	Discrete	×	—	OSE/NE	—
[11]	Stack. Hyper.	Generic	✓	×	SHE	✓
[16]	Stack.+Match.	T+VQA	✓	×	SE+Stab.	Part.
[17]	Wt. Potential	Text	✓	Mod.	NE	✓
[19]	MF-Stack.	Power	✓	✓	ϵ -NE	✓
Proposed	MF-Stack.	T+I+S	✓	✓	ϵ-NE	✓

MU: multiuser; Modal.: modality; Scal.: scalable; Equil.: equilibrium; Dist.: distributed; Mod.: moderate; Part.: partial; Wt.: weighted; Seq.: sequential; Stack.: Stackelberg; Hyper.: hypergame; Match.: matching; Stab.: stability; MFE: mean-field equilibrium; T/I/S: text/image/speech; PHY: physical layer; VQA: visual question answering.

tiagent reinforcement learning to allocate vehicular network resources. However, these studies have not addressed SemCom or modality-specific quality functions with per-modality congestion coupling that characterize multimodal SemCom.

B. DRL and Other Approaches

Resource allocation in semantic-aware networks has garnered increased attention [24]–[26], with DRL emerging as a prominent approach. Particularly, Hu *et al.* [27] and Chen *et al.* [3] applied DRL to allocating multimodal SemCom resources, performing effectively but requiring centralized training with global channel state information and the full knowledge of all users' semantic preferences, incurring prohibitive signaling overhead as the network scales.

Optimization-based methods are analytically tractable but have limited scope. For instance, Yan *et al.* [4] defined semantic spectral efficiency for text SemCom. In addition, Xia *et al.* [28] optimized hybrid semantic/bit networks using convex relaxation. Further, Xu *et al.* [29] developed semantic-aware power allocation for generative image SemCom, and Hua *et al.* [30] analyzed tradeoffs between sensing and transmission bandwidth. These studies are all restricted to single-modality settings and depend on centralized computation, which becomes impractical as the network scales and lacks equilibrium guarantees for networks with autonomous agents.

Positioning of This Work

Table I compares the proposed framework with representative studies. The closest SemCom Stackelberg work uses a uniform price and DRL-based allocation [6], while existing MFG-based baselines focus on conventional power, computing, or binary semantic decisions [7], [22]. Accordingly, the proposed framework is positioned around three combined capabilities: continuous per-modality bandwidth allocation with calibrated modality-specific quality functions, an ϵ -Nash equilibrium guarantee with an $O(1/\sqrt{N})$ approximation rate, and equilibrium computation whose dominant fixed-point step is independent of the number of users.

III. SYSTEM MODEL AND PROBLEM FORMULATION

In this work, we consider a single-cell downlink SemCom network where a base station (BS) serves N users, each

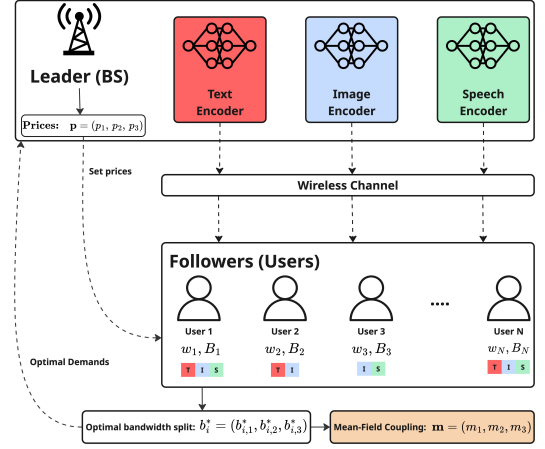


Fig. 1: System model of the mean-field Stackelberg equilibrium (MFSE) framework to allocate multimodal semantic communication (SemCom) resources ($K = 3$ modalities).

requiring multimodal semantic content delivery across text ($k = 1$), image ($k = 2$), and speech ($k = 3$) [1], [2]. The BS operates under a fixed total system bandwidth B_{total} , with $B_{\text{total}} = 20$ MHz used as the primary scalability setting in the experiments. Accordingly, admission control and network slicing induce the nominal per-user budget $B_0 = B_{\text{total}}/N$. To handle the joint challenge of allocating the limited bandwidth across modalities and users while accounting for heterogeneous task requirements, time-varying channels, and diverse resource–quality relationships per modality, we formulate a hierarchical game-theoretic framework as follows.

As depicted in Fig. 1, the BS acts as a Stackelberg leader with three modality-specific neural encoders that transmit semantic content to N users over wireless channels. The BS sets a modality pricing vector $\mathbf{p} = (p_1, p_2, p_3)$ and broadcasts it to all users. Each user (Stackelberg follower) is characterized by a task weight w_i and bandwidth budget B_i and computes an optimal bandwidth split $\mathbf{b}_i^* = (b_{i,1}^*, b_{i,2}^*, b_{i,3}^*)$ across the three modalities. These individual demands are aggregated into a mean-field coupling vector $\mathbf{m} = (m_1, m_2, m_3)$, capturing the population-average demand per modality. In this hierarchical structure, the leader sets prices, and the followers respond with optimal demands, forming the basis of the MFSE framework detailed below. Table II summarizes the main notation used in the model.

A. Network Architecture and Channel Model

The BS uses modality-specific semantic encoders that compress source information into task-relevant representations. Let $\mathbf{w}_i = (w_{i,1}, w_{i,2}, w_{i,3}) \in \mathbb{R}_+^3$ denote the weight vector of user $i \in \{1, \dots, N\}$, drawn independently from $\mathcal{F}_w$, where $w_{i,k}$ represents the importance of modality k to the task of user i . In addition, communication occurs over Rayleigh fading channels with channel power gains that are independent across users and modalities, which is given by

$$h_{i,k} \sim \text{Exp}(\sigma_h^2), \quad \forall i \in \{1, \dots, N\}, \quad k \in \{1, 2, 3\}. \quad (1)$$

Let P_0 denote the per-user transmission power budget, with aggregate BS transmit power $P_{\text{tx}} = N \cdot P_0$, and let σ_n^2 denote

TABLE II: Notation summary.

Symbol	Meaning
i	User index, $i \in \{1, \dots, N\}$.
k	Modality index, $k \in \{1, 2, 3\}$ for text, image, and speech.
B_{total}	Total system bandwidth.
B_0	Nominal per-user bandwidth, $B_0 = B_{\text{total}}/N$ in the primary fixed-total regime.
B_i	Bandwidth budget assigned to user i .
$b_{i,k}, \mathbf{b}_i$	Bandwidth allocated by user i to modality k and the corresponding allocation vector.
$h_{i,k}, \mathbf{h}_i$	Channel power gain of user i on modality k and the corresponding channel vector.
$w_{i,k}, \mathbf{w}_i$	Task-importance weight of user i on modality k and the corresponding weight vector.
$a_k, \alpha_k, \kappa_k, \delta, p$	Modality-specific quality parameters calibrated to real codecs (Section V-B).
γ	Dimensionless channel-rate scaling parameter.
$\gamma_{\min}$	Minimum operating threshold for the concave codec regime.
$p_k, \mathbf{p}$	Price of modality k and the modality pricing vector.
$m_k, \mathbf{m}$	Mean-field demand of modality k and the mean-field vector.
$\eta_k, \boldsymbol{\eta}$	Congestion coefficient of modality k and the congestion vector.
ρ	Welfare weight in the BS objective.
λ, μ_k	Budget and floor Lagrange multipliers used in Proposition 1.
μ_k^{svc}	Service rate of the shared semantic encoder for modality k .

the noise power spectral density. In this work, we define the channel-rate scaling parameter $\gamma = P_0/(\sigma_n^2 \cdot B_{\text{ref}})$, where B_{ref} represents a reference bandwidth (set to B_0 without loss of generality) that renders γ dimensionless. Accordingly, the quality functions depend on the product $b_{i,k} \cdot h_{i,k} \cdot \gamma$, which represents the resource-channel quality for user i on modality k . The modality-specific rate coefficients $(\kappa_1, \kappa_2, \delta)$ absorb the B_{ref} normalization; hence, the quality-function arguments are dimensionless.

B. Modality-Specific Semantic Quality Functions

Without loss of generality, we adopt modality-specific quality functions to capture the distinct resource–quality relationship of each modality stream [2]. Each function is calibrated to a downstream task metric of a real neural codec: image top-1 classification accuracy (Deep JSCC on CIFAR-10), text BLEU (DeepSC on Europarl), and speech PESQ (DeepSC-S on VoiceBank-DEMAND). Because these measured metrics are bounded and saturate with the transmission rate, text and image quality are modeled by a bounded capacity-compressed surrogate and speech intelligibility by a saturating power-compressed form. Accordingly, the modality-specific quality functions can be expressed as

$$Q_{\text{text}}(b, h) = a_1 \left[1 - (1 + \kappa_1 b h \gamma)^{-\alpha_1} \right], \quad (2)$$

$$Q_{\text{image}}(b, h) = a_2 \left[1 - (1 + \kappa_2 b h \gamma)^{-\alpha_2} \right], \quad (3)$$

$$Q_{\text{speech}}(b, h) = a_3 \tanh(\delta (b h \gamma)^p), \quad (4)$$

where $a_k > 0$ denotes the modality-specific saturation ceiling and $(\alpha_k, \kappa_k, \delta, p) > 0$ are the remaining calibrated shape parameters. These parameters are fitted to the real-codec measurements in Section V-B, where the fitted values and goodness-of-fit are reported. For simplified notation, $Q_k(b, h)$ denotes the quality function for modality k . Over the validated operating region (Assumption 1), each Q_k is bounded above by a_k , continuously differentiable, strictly increasing, and strictly concave in b for $h > 0$, with $Q_k(0, h) = 0$. All subsequent analysis relies only on this generic bounded, increasing,

and strictly concave property, rather than on the superseded unbounded exponential/logarithmic/hyperbolic-tangent forms. The calibrated text and image forms are strictly concave for all $b \geq 0$, while the QoS floors $b_k^{\min} \geq 0$ in (21) (positive for the image modality) together with the activation conditions in (27)–(28) confine every active allocation to the increasing, strictly concave branch, so the knee restriction is load-bearing only for the power-compressed speech modality.

Assumption 1 (Operating Regime). *The network operates above a modality-specific concavity knee $\gamma_{\min} > 0$ such that the resource-channel quality satisfies $b_{i,k} \cdot h_{i,k} \cdot \gamma \geq \gamma_{\min}$ for all active allocations. Real neural codecs exhibit an abrupt, convex degradation at very low channel quality; above the knee, the measured task-metric curves are monotone and concave, so the calibrated models (2) to (4) are used only in this validated regime. The fitted knees obtained from the real-codec calibration in Section V-B are $\gamma_{\min} \approx 9.86$ dB for text, 9.94 dB for image, and 13.01 dB for speech. The KKT conditions in Proposition 1 require strict concavity, which holds throughout this region. This assumption is mild because the evaluation operates at a system SNR of 20 dB (Section V), comfortably above all three knees [1]. Because the calibrated text and image models are strictly concave for all $b \geq 0$, the knee restriction governs the fidelity of the fits to the measured codec curves and the concavity of the power-compressed speech branch, rather than the validity of the KKT development.*

C. Resource Allocation and Constraints

Under the primary fixed-total-bandwidth model, the nominal budget of each admitted user is induced by the total system bandwidth, which is given by

$$B_i = B_0 = \frac{B_{\text{total}}}{N}, \quad \forall i \in \{1, \dots, N\}. \quad (5)$$

This framing keeps B_{total} fixed as the number of users changes, so the per-user nominal budget decreases with N . Unless explicitly stated otherwise, all scalability statements in this paper use this fixed-total convention. A constant-per-user setting with B_0 held fixed as N varies is treated only as an ablation or legacy comparison, in which the total bandwidth would scale as $B_{\text{total}}(N) = NB_0$.

Remark 1 (Fixed-Total Provisioning). *For any fixed N , the value of $B_0 = B_{\text{total}}/N$ is determined by admission control before the follower game is solved. Therefore, each user's feasible set $\mathcal{B}_i$ remains independent of the instantaneous decisions of other users, as required by the mean-field independence in (14). The preslicing concession is that users cannot share bandwidth across the per-user partition. A global constraint $\sum_{i,k} b_{i,k} \leq B_{\text{total}}$ would enable multiuser diversity but compromise the MFG structure.*

Accordingly, each user i allocates this budget across modalities, with $\mathbf{b}_i = (b_{i,1}, b_{i,2}, b_{i,3}) \in \mathbb{R}_+^3$ subject to the following:

$$\sum_{k=1}^3 b_{i,k} \leq B_i, \quad b_{i,k} \geq 0, \quad \forall k \in \{1, 2, 3\}. \quad (6)$$

D. User Utility Function

Following the game-theoretic framework in [5], the utility of user i balances the semantic quality against the resource costs. Accordingly, the utility function is formulated as

$$U_i(\mathbf{b}_i; \mathbf{p}, \mathbf{m}) = \sum_{k=1}^3 w_{i,k} Q_k(b_{i,k}, h_{i,k}) - \sum_{k=1}^3 p_k b_{i,k} - \sum_{k=1}^3 \eta_k b_{i,k} m_k, \quad (7)$$

where $\mathbf{p} = (p_1, p_2, p_3) \in \mathbb{R}_+^3$ represents the modality-specific pricing vector, $\mathbf{m} = (m_1, m_2, m_3) \in \mathbb{R}_+^3$ indicates the mean-field vector representing the population-average bandwidth demand, and $\boldsymbol{\eta} = (\eta_1, \eta_2, \eta_3) \in \mathbb{R}_{++}^3$ denotes the congestion-coupling coefficients.

Remark 2 (Semantic-Priority Decomposition). *The task-importance weight $w_{i,k}$ can be decomposed into a base preference component $\bar{w}_{i,k} > 0$ and a semantic-priority multiplier $\pi_{i,k} \geq 0$ that reflects task-criticality or content-importance, i.e., $w_{i,k} = \bar{w}_{i,k} \pi_{i,k}$. Under this decomposition the utility (7) becomes*

$$U_i = \sum_{k=1}^3 \bar{w}_{i,k} \pi_{i,k} Q_k(b_{i,k}, h_{i,k}) - \sum_{k=1}^3 (p_k + \eta_k m_k) b_{i,k}, \quad (8)$$

and the KKT analysis in Proposition 1 carries through unchanged with the effective weight $w_{i,k} = \bar{w}_{i,k} \pi_{i,k}$. Heterogeneous priority profiles π_i steer allocations toward modalities with higher task-criticality, as demonstrated experimentally in Section V.

In particular, the three terms capture weighted semantic quality (WSQ), per-modality pricing payments, and computational congestion at the BS's shared semantic processing pipeline. The congestion term $\eta_k b_{i,k} m_k$ models the queueing delay at the BS's GPU-accelerated encoder for modality k , such as DeepSC for text, joint source-channel coding (JSCC) for image, and DeepSC-S for speech. It is worth noting that all users share these neural encoders; hence, the aggregate demand reflected by $m_k = \mathbb{E}[b_k^*]$ determines the computational load on the k -th encoder [7]. This computational congestion arises from shared neural processing resources, not radio frequency (RF) interference, because the bandwidth of each user is orthogonalized. To justify the linear coupling, consider an M/M/1 encoder with service rate μ_k^{svc} and normalized load m_k . The excess delay satisfies

$$D_k(m_k) = \frac{m_k}{\mu_k^{\text{svc}}(\mu_k^{\text{svc}} - m_k)} = \frac{m_k}{(\mu_k^{\text{svc}})^2} + O(m_k^2), \quad (9)$$

for $m_k \ll \mu_k^{\text{svc}}$. Therefore, the first-order delay cost coupled with the user's own demand is $\eta_k b_{i,k} m_k$ with $\eta_k = 1/(\mu_k^{\text{svc}})^2$. Near saturation, the exact delay diverges and the linear approximation underestimates congestion; this nonlinear regime is outside the light-to-moderate traffic setting considered by the closed-form KKT development.

E. Stackelberg Game Formulation

To handle the multiple-target objective of maximizing the semantic quality while maintaining the computational efficiency, we formulate the BS-user interaction as a Stackelberg game [6].

1) *Follower Problem (Users)*: Given price $\mathbf{p}$ and mean-field vector $\mathbf{m}$, each user i solves the following:

$$\max_{\mathbf{b}_i \in \mathcal{B}_i} \sum_{k=1}^3 w_{i,k} Q_k(b_{i,k}, h_{i,k}) - \sum_{k=1}^3 (p_k + \eta_k m_k) b_{i,k}, \quad (10)$$

where $\mathcal{B}_i = \{\mathbf{b} \in \mathbb{R}_+^3 : \sum_k b_k \leq B_i\}$. Each Q_k is strictly concave, and the costs are linear. Hence, the problem admits a unique optimal allocation $\mathbf{b}_i^*(\mathbf{p}, \mathbf{m}, \mathbf{h}_i, \mathbf{w}_i)$.

2) *Leader Problem (Base Station)*: To optimize the tradeoff between revenue and social welfare, the BS solves the following:

$$\max_{\mathbf{p} \in [\underline{p}, \bar{p}]^3} \Phi(\mathbf{p}) = (1 - \rho) R(\mathbf{p}) + \rho W(\mathbf{p}), \quad (11)$$

where $\rho \in [0, 1]$ denotes the welfare weight, and the social welfare is given by

$$W(\mathbf{p}) = N \cdot \mathbb{E}_{\mathbf{h}, \mathbf{w}} \left[\sum_{k=1}^3 w_k Q_k(b_k^*(\mathbf{p}, \mathbf{m}^*(\mathbf{p}), \mathbf{h}, \mathbf{w}), h_k) \right]. \quad (12)$$

The revenue aggregates per-user payments using the mean-field consistency. This value is computed over only the price-responsive portion of demand to prevent the BS from profiting from floor-forced allocation (see Remark 3), which is expressed as

$$R(\mathbf{p}) = N \sum_{k=1}^3 p_k (m_k^*(\mathbf{p}) - b_k^{\min})^+, \quad (13)$$

where $(\cdot)^+ = \max(\cdot, 0)$. Moreover, $b_k^{\min}$ denotes the min-QoS floor from (21), $\mathbf{m}^*(\mathbf{p})$ indicates the equilibrium mean-field at price $\mathbf{p}$, and the parameters $(\underline{p}, \bar{p})$ bound the admissible prices. The parameter ρ enables a tunable tradeoff between the operator revenue ($\rho = 0$) and user experience ($\rho = 1$). When $b_k^{\min} = 0$, the revenue reduces to $R(\mathbf{p}) = N \sum_k p_k m_k^*(\mathbf{p})$.

F. Mean-Field Approximation

For large N , the N -player game is intractable. Therefore, following the mean-field game (MFG) framework [7], we replace finite-player interactions with a mean-field approximation. In the finite game, the congestion cost is $C_{i,k}^{\text{cong}} = \eta_k b_{i,k} \cdot \frac{1}{N} \sum_{j=1}^N b_{j,k}$. As $N \rightarrow \infty$, the law of large numbers yields the following:

$$\frac{1}{N} \sum_{j=1}^N b_{j,k}^* \xrightarrow{N \rightarrow \infty} \mathbb{E}_{\mathbf{h}, \mathbf{w}} [b_k^*(\mathbf{p}, \mathbf{m}, \mathbf{h}, \mathbf{w})] \triangleq m_k, \quad (14)$$

which motivates the MFG formulation where each user interacts with aggregate behavior $\mathbf{m} = (m_1, m_2, m_3)$. Under the primary fixed-total model, $B_0 = B_{\text{total}}/N$ is fixed once the population size is specified, and each user's optimal allocation remains independent of any other individual user's action. When scalability is studied by varying N , m_k is interpreted

under the corresponding induced budget $B_0(N)$, while any constant-per-user asymptotic regime is explicitly labelled as a separate ablation. Accordingly, the mean-field must satisfy the fixed-point consistency condition, which can be expressed as

$$m_k = \mathbb{E}_{\mathbf{h}, \mathbf{w}}[b_k^*(\mathbf{p}, \mathbf{m}, \mathbf{h}, \mathbf{w})], \quad k = 1, 2, 3. \quad (15)$$

Definition 1 (Mean-Field Stackelberg Equilibrium). *A tuple $(\mathbf{p}^*, \mathbf{m}^*)$ constitutes an MFSE if the following conditions are met:*

- 1) **Mean-field consistency:** For each $k \in \{1, 2, 3\}$,

$$m_k^* = \mathbb{E}_{\mathbf{h}, \mathbf{w}}[b_k^*(\mathbf{p}^*, \mathbf{m}^*, \mathbf{h}, \mathbf{w})]; \quad (16)$$

- 2) **User optimality:** For almost every $(\mathbf{h}, \mathbf{w})$,

$$\mathbf{b}^*(\mathbf{p}^*, \mathbf{m}^*, \mathbf{h}, \mathbf{w}) = \arg \max_{\mathbf{b} \in \mathcal{B}} U(\mathbf{b}; \mathbf{p}^*, \mathbf{m}^*, \mathbf{h}, \mathbf{w}); \quad (17)$$

- 3) **BS optimality:** $\mathbf{p}^* = \arg \max_{\mathbf{p} \in [\underline{p}, \bar{p}]^3} \Phi(\mathbf{p}, \mathbf{m}^*(\mathbf{p}))$.

The MFSE captures a hierarchical equilibrium, where the BS anticipates the mean-field response to its pricing, whereas each user optimizes against aggregate behavior rather than tracking individual competitors. As a result, this approach decouples equilibrium computation from N , while deployment over realized users scales with the number of admitted users, thereby yielding the scalable framework detailed in Section IV.

G. Cross-Modal Correlation and Shared-Feature Extension

To model cross-modal semantic correlation (e.g., an image and its caption, or speech and its transcript), we introduce a shared-feature coupling at the resource-allocation level. Let $C \in [0, 1]^{K \times K}$ be a symmetric, hollow correlation matrix whose entry C_{kj} is the fraction of modality k 's semantic content that is also conveyed by a correlated modality j , and let $\rho_c \in [0, 1]$ be a global correlation-strength dial. A correlated modality j then donates redundant semantic resource to modality k , so the *effective* bandwidth that drives semantic quality is

$$b_{i,k}^{\text{eff}} = b_{i,k} + \rho_c \sum_{j \neq k} C_{kj} b_{i,j}, \quad \mathbf{b}_i^{\text{eff}} = R(\rho_c) \mathbf{b}_i, \quad R(\rho_c) = I + \rho_c C \quad (18)$$

so that $\rho_c = 0$ recovers the independent model ($R = I$). Because $b_{i,k}^{\text{eff}} \geq b_{i,k}$, shared features let a modality reach a target quality with less of its own spectrum; the freed native bandwidth is the reported cross-modal saving, while quality is evaluated on the effective bandwidth $Q_k(b_{i,k}^{\text{eff}}, h_{i,k})$. The per-user stationarity condition then carries a cross-modal coupling term, $\sum_l w_{i,l} Q'_l((R\mathbf{b}_i)_l, h_{i,l}) R_{lk} = p_k + \eta_k m_k + \lambda$ for each active modality k . Because R is constant in $\mathbf{b}$, the per-user Hessian $R^\top \text{diag}(w_{i,l} Q''_l) R$ remains negative definite whenever $\rho_c \|C\| < 1$, so the follower problem stays strictly concave with a unique optimum, and $\rho_c = 0$ collapses the condition to the decoupled KKT form of Proposition 1; the mean-field fixed point (15) is otherwise unchanged. The correlation matrix used in the evaluation is $C_{\text{text,image}} = 0.60$, $C_{\text{text,speech}} = 0.40$, and $C_{\text{image,speech}} = 0.25$, and the impact of ρ_c is analyzed in Section V.

IV. PROPOSED SOLUTION

In this section, we develop an efficient algorithm to compute the MFSE defined in Definition 1. To address the hierarchical structure of the problem, we design the solution procedure with two nested components: (i) an inner MFG fixed-point iteration that computes the equilibrium $\mathbf{m}^*$ for the given prices $(\mathbf{p})$, and (ii) an outer Stackelberg loop, where the BS updates prices toward a stationary/KKT point using the MFG equilibrium mapping. Accordingly, this section derives the analytical structure of the user optimization problem, presents the complete algorithm, and establishes its theoretical properties.

A. KKT-Based User Optimization

Given prices $\mathbf{p}$, mean-field vector $\mathbf{m}$, channel realization $\mathbf{h}$, and weights $\mathbf{w}$, each user solves

$$\max_{\mathbf{b} \in \mathcal{B}} \sum_{k=1}^3 w_k Q_k(b_k, h_k) - \sum_{k=1}^3 c_k b_k, \quad (19)$$

where $c_k \triangleq p_k + \eta_k m_k$ represents the effective marginal cost of the bandwidth on modality k . To model the congestion at the shared modality- k encoder, we adopt a light-load approximation of an M/M/1 queueing system. Specifically, let μ_k^{svc} denote the normalized service rate of the modality- k encoder under aggregate offered load m_k . The average waiting delay $W_k(m_k) = m_k / (\mu_k^{\text{svc}} (\mu_k^{\text{svc}} - m_k))$ can be expanded around $m_k = 0$ as $W_k(m_k) = m_k / (\mu_k^{\text{svc}})^2 + O(m_k^2)$. Accordingly, the user-level congestion cost is calculated as:

$$C_k^{\text{cong}}(b_k, m_k) = b_k W_k(m_k) \simeq \eta_k b_k m_k, \quad \eta_k = \frac{1}{(\mu_k^{\text{svc}})^2}. \quad (20)$$

This linear congestion cost serves as a first-order approximation of the queueing delay in the light-to-moderate load regime. Hence, it preserves the convexity and analytical tractability of the KKT conditions, while the exact M/M/1 delay or more complex GPU scheduling models are left for high-load scenarios. Let $\mathcal{B}$ denote the feasible set incorporating per-modality minimum quality-of-service (QoS) floors, which is expressed as

$$\mathcal{B} = \left\{ \mathbf{b} \in \mathbb{R}_+^3 : \sum_k b_k \leq B_0, b_k \geq b_k^{\min}, \forall k \right\}, \quad (21)$$

with $b_k^{\min} \geq 0$ indicating the minimum bandwidth guarantee for modality k . Setting $b_k^{\min} = 0$ recovers the unconstrained case. The floor prevents resource starvation for modalities whose calibrated quality functions saturate at different rates under the real-codec operating curves. Each deployed Q_k is increasing and strictly concave on the active calibrated branch, and $\mathcal{B}$ denotes a compact convex polytope. Hence, the KKT conditions are necessary and sufficient for the user-level optimum over that operating branch.

Proposition 1 (Optimal User Allocation). *The unique optimal bandwidth allocation, denoted by $\mathbf{b}^* = (b_1^*, b_2^*, b_3^*)$, satisfies one of three scenarios for each modality:*

$$w_k \cdot Q'_k(b_k^*, h_k) = c_k + \lambda^* - \nu_k, \quad \forall k, \quad (22)$$

with complementary slackness denoted by $\nu_k(b_k^* - b_k^{\min}) = 0$ and $\nu_k \geq 0$, where $\lambda^* \geq 0$ indicates the Lagrange multiplier for the budget constraint $\sum_k b_k \leq B_0$, and ν_k denotes the multiplier for the min-QoS floor $b_k \geq b_k^{\min}$. Three regimes arise: (i) $b_k^* > b_k^{\min}$ (unconstrained, $\nu_k = 0$), where the stationarity reduces to $w_k Q'_k = c_k + \lambda^*$; (ii) $b_k^* = b_k^{\min}$ (floor-binding, $\nu_k > 0$), where allocation is pinned at the minimum guarantee; and (iii) $b_k^* = b_k^{\min} = 0$ (inactive), where $Q'_k(b, h) = \partial Q_k(b, h)/\partial b$ denotes the marginal semantic quality.

Proof. The Lagrangian of (19) subject to (21) is given by

$$\mathcal{L} = \sum_k w_k Q_k(b_k, h_k) - \sum_k c_k b_k - \lambda \left(\sum_k b_k - B_0 \right) + \sum_k \nu_k (b_k - b_k^{\min}). \quad (23)$$

The KKT stationarity condition $\partial \mathcal{L} / \partial b_k = 0$ results in (22). Complementary slackness $\nu_k(b_k - b_k^{\min}) = 0$ with $\nu_k \geq 0$ yields three regimes. Uniqueness follows from the strict concavity of the objective over the convex feasible set (21). This completes the proof. $\square$

To obtain the per-user best response, we invert the KKT condition (22) for each modality. Computing the derivatives of the calibrated capacity-compressed forms (2)–(4) results in the following:

$$Q'_{\text{text}}(b, h) = a_1 \alpha_1 \kappa_1 h \gamma (1 + \kappa_1 b h \gamma)^{-(\alpha_1+1)}, \quad (24)$$

$$Q'_{\text{image}}(b, h) = a_2 \alpha_2 \kappa_2 h \gamma (1 + \kappa_2 b h \gamma)^{-(\alpha_2+1)}, \quad (25)$$

$$Q'_{\text{speech}}(b, h) = a_3 \delta p h \gamma (b h \gamma)^{p-1} \text{sech}^2(\delta (b h \gamma)^p). \quad (26)$$

By setting $w_k Q'_k(b_k^*, h_k) = c_k + \lambda$ and solving for b_k^* , the text and image modalities admit closed-form allocations, whereas the power-compressed speech modality is solved by a one-dimensional monotone bisection:

$$b_k^* = \left[\frac{1}{\kappa_k \tilde{h}_k} \left(\left(\frac{w_k a_k \alpha_k \kappa_k \tilde{h}_k}{c_k + \lambda} \right)^{\frac{1}{\alpha_k+1}} - 1 \right) \right]^+, \quad k \in \{1, 2\}, \quad (27)$$

$$b_3^* : w_3 Q'_3(b_3^*, h_3) = c_3 + \lambda, \quad b_3^* \in [0, B_0], \quad (28)$$

where $\tilde{h}_k \triangleq h_k \gamma$ denotes the channel quality. The text and image branches are active only when the marginal utility at the origin exceeds the effective cost, i.e., $w_k a_k \alpha_k \kappa_k \tilde{h}_k > c_k + \lambda$; otherwise the closed form returns zero before the floor is applied. For the power-compressed speech modality with $p < 1$, the marginal utility satisfies $Q'_3(0^+, h) \rightarrow \infty$, so (28) always admits a unique positive root; because Q'_3 is strictly decreasing in b , this root is obtained by monotone bisection within $O(\log(1/\varepsilon))$ iterations. The zero or sub-floor branch is then replaced with $b_k^{\min}$ to enforce the min-QoS floor. When the floor is binding ($b_k^* = b_k^{\min}$), the allocation is fixed regardless of price, and the corresponding sensitivity is $\partial b_k^* / \partial c_j = 0$ (see Lemma 1).

1) Bisection on the Lagrange Multiplier: To determine the Lagrange multiplier λ^* , we employ the budget constraint. If the unconstrained solution (setting $\lambda = 0$) satisfies $\sum_k b_k^*(0) \leq B_0$, then $\lambda^* = 0$, and the constraint is inactive. Otherwise, λ^* represents the unique root of the following:

$$g(\lambda) \triangleq \sum_{k=1}^3 b_k^*(\lambda) - B_0 = 0, \quad (29)$$

which can be determined using bisection. Each $b_k^*(\lambda)$ is monotonically decreasing in λ ; hence, bisection converges geometrically within $O(\log(1/\varepsilon))$ iterations.

B. Mean-Field Game Fixed-Point Iteration

For the given prices ($\mathbf{p}$), the MFG equilibrium $\mathbf{m}^*(\mathbf{p})$ is a fixed point of the mapping

$$T(\mathbf{m}) = \mathbb{E}_{\mathbf{h}, \mathbf{w}}[\mathbf{b}^*(\mathbf{p}, \mathbf{m}, \mathbf{h}, \mathbf{w})]. \quad (30)$$

Proposition 2 (Mean-Field Mapping Convergence). *Suppose the channel gains and weights have a continuous joint density on $\mathbb{R}_+^3 \times \mathbb{R}_+^3$ with finite first moments, which is satisfied by the Rayleigh (exponential) channel model (1). Then, under the condition*

$$\|\bar{\mathbf{J}} \text{diag}(\boldsymbol{\eta})\|_{\infty} < 1, \quad (31)$$

where $\bar{\mathbf{J}} = \mathbb{E}_{\mathbf{h}, \mathbf{w}}[\mathbf{J}(\mathbf{h}, \mathbf{w})]$ denotes the population-averaged per-user Jacobian from Lemma 1, the mapping T in (30) represents a contraction in the ℓ^{∞} norm on the domain $[0, B_0]^3$.

Proof. We consider two mean-field vectors $\mathbf{m}, \mathbf{m}' \in [0, B_0]^3$ with the following cost vectors:

$$\mathbf{c} = \mathbf{p} + \text{diag}(\boldsymbol{\eta}) \mathbf{m}, \quad \mathbf{c}' = \mathbf{p} + \text{diag}(\boldsymbol{\eta}) \mathbf{m}'. \quad (32)$$

For a given user with type $(\mathbf{h}, \mathbf{w})$ and optimal allocation $\mathbf{b}^*(\mathbf{c})$, let the per-user Jacobian be defined as

$$J_{kj} \triangleq \frac{\partial b_k^*}{\partial c_j}, \quad (33)$$

accounting for the direct effect of cost changes and the cross-modality coupling via λ^* when the budget constraint is active (Lemma 1). The mean value theorem applied componentwise results in

$$\begin{aligned} |T_k(\mathbf{m}) - T_k(\mathbf{m}')| &= \left| \mathbb{E} \left[\sum_{j=1}^3 J_{kj} \eta_j (m_j - m'_j) \right] \right| \\ &\leq \sum_{j=1}^3 |\bar{J}_{kj}| \eta_j |m_j - m'_j|. \end{aligned} \quad (34)$$

Taking the ℓ^{∞} norm over k yields the following:

$$\|T(\mathbf{m}) - T(\mathbf{m}')\|_{\infty} \leq \|\bar{\mathbf{J}} \text{diag}(\boldsymbol{\eta})\|_{\infty} \|\mathbf{m} - \mathbf{m}'\|_{\infty}, \quad (35)$$

where $\|\cdot\|_{\infty}$ represents the induced matrix infinity norm (maximum absolute row sum). When (31) holds, T denotes a contraction, and the Banach fixed-point theorem ensures the existence of a unique fixed point $\mathbf{m}^*$.

Next, we verify that $\bar{\mathbf{J}} = \mathbb{E}[\mathbf{J}]$ has finite entries. For text and images, Q'_k in (41) is continuous and strictly negative

on the active region, so D_k is bounded away from zero and the Jacobian entries are uniformly bounded. For the power-compressed speech modality with $p < 1$, $Q'_3(0^+, h) \rightarrow \infty$ and $Q''_3(0^+, h) \rightarrow -\infty$, so the active allocation lies in the strictly concave interior and $J_{33} = 1/D_3 \rightarrow 0$ as $b_3^* \rightarrow 0^+$; hence no boundary singularity arises and $|J_{33}|$ is bounded on the active region. Because these bounds hold on the active region irrespective of the channel magnitude, $\bar{J} = \mathbb{E}[\mathbf{J}]$ is finite even under the unbounded exponential channel law (1). When the contraction condition is not met, Brouwer's theorem still guarantees at least one fixed point. This completes the proof. $\square$

In practice, a damped update is employed to improve robustness beyond the sufficient condition in (31), which is given by

$$\mathbf{m}^{(t+1)} = (1 - \theta) \mathbf{m}^{(t)} + \theta T(\mathbf{m}^{(t)}), \quad (36)$$

where $\theta \in (0, 1]$ denotes the damping factor. Setting $\theta = 1$ recovers the standard Picard iteration, whereas smaller values of θ improve stability when the raw infinity-norm contraction condition is violated.

Proposition 3 (Damped Mean-Field Contraction). *Let $\mathbf{J}_T \triangleq \bar{J} \text{diag}(\boldsymbol{\eta})$ be the Jacobian of T at a fixed point, and define the damped map*

$$T_\theta(\mathbf{m}) \triangleq (1 - \theta)\mathbf{m} + \theta T(\mathbf{m}). \quad (37)$$

If $s \triangleq \max_i \text{Re}(\lambda_i(\mathbf{J}_T)) < 1$, then there exists $\theta^ \in (0, 1]$ such that T_θ is a contraction in a weighted norm for every $\theta \in (0, \theta^*]$.*

Proof. The Jacobian of T_θ is $\mathbf{G}_\theta = (1 - \theta)\mathbf{I} + \theta\mathbf{J}_T$, yielding eigenvalues $g_i(\theta) = 1 - \theta(1 - \lambda_i(\mathbf{J}_T))$. Let $1 - \lambda_i(\mathbf{J}_T) = a_i - id_i$, with $a_i = 1 - \text{Re}(\lambda_i(\mathbf{J}_T)) > 0$. The squared modulus $|g_i(\theta)|^2 = (1 - \theta a_i)^2 + \theta^2 d_i^2$ has derivative $-2a_i < 0$ at $\theta = 0$. Thus, $|g_i(\theta)| < 1$ for $\theta \in (0, \theta^*]$ with $\theta^* = \min_i 2a_i/(a_i^2 + d_i^2)$. The spectral radius of $\mathbf{G}_\theta$ is strictly less than unity, ensuring contraction. This completes the proof. $\square$

Proposition 3 relaxes (31) to only requiring real parts of eigenvalues of $\mathbf{J}_T$ to be less than one. For instance, at 10 dB SNR where the undamped constant is 4.39, the damping $\theta = 0.3$ stabilizes the fixed point with a spectral radius of $0.70 < 1$. At higher SNRs, the undamped constant is small (0.004–0.006), ensuring convergence.

The expectation is approximated using N_{mc} Monte Carlo (MC) samples, which is calculated as

$$\hat{T}(\mathbf{m}) = \frac{1}{N_{\text{mc}}} \sum_{s=1}^{N_{\text{mc}}} \mathbf{b}^*(\mathbf{p}, \mathbf{m}, \mathbf{h}^{(s)}, \mathbf{w}^{(s)}). \quad (38)$$

The optimal allocation $\mathbf{b}^*(\mathbf{c})$ features rectified truncation at the active-set boundary, making $\hat{T}(\mathbf{m})$ only piecewise smooth. However, the probability of sampling exactly at a nondifferentiable kink is measure zero under continuous $(\mathbf{h}, \mathbf{w})$ distributions, and the Adam optimizer's moment averaging further smooths residual nondifferentiability.

Algorithm 1 summarizes the inner MFG iteration.

Algorithm 1 Mean-Field Game Fixed-Point Iteration

Require: Price $\mathbf{p}$, MC samples $\{(\mathbf{h}^{(s)}, \mathbf{w}^{(s)})\}_{s=1}^{N_{\text{mc}}}$, damping θ , tolerance ε_{mfg}

Ensure: Equilibrium mean-field vector $\mathbf{m}^*$

```

1: Initialize  $\mathbf{m}^{(0)} = (B_0/3, B_0/3, B_0/3)$ 
2: for  $t = 0, 1, 2, \dots$  do
3:   Compute effective costs:  $c_k \leftarrow p_k + \eta_k m_k^{(t)}, \forall k$ 
4:   for  $s = 1, \dots, N_{\text{mc}}$  (in parallel) do
5:     Solve (27)–(28) using analytic text/image inverses,
       monotone speech bisection, and bisection on  $\lambda$  to
       obtain  $\mathbf{b}^{*(s)}$ 
6:   end for
7:   Compute  $\hat{T}(\mathbf{m}^{(t)}) = \frac{1}{N_{\text{mc}}} \sum_s \mathbf{b}^{*(s)}$ 
8:   Damped update:  $\mathbf{m}^{(t+1)} \leftarrow (1 - \theta) \mathbf{m}^{(t)} + \theta \hat{T}(\mathbf{m}^{(t)})$ 
9:   if  $\|\mathbf{m}^{(t+1)} - \mathbf{m}^{(t)}\| < \varepsilon_{\text{mfg}}$  then
10:    return  $\mathbf{m}^* \leftarrow \mathbf{m}^{(t+1)}$ 
11:  end if
12: end for
```

C. Stackelberg Pricing via Implicit Differentiation

At the outer level, the BS updates the price vector $\mathbf{p}$ toward a stationary/KKT point of the nonconvex bilevel problem in (11). The BS objective $\Phi(\mathbf{p})$ directly depends on prices and the MFG equilibrium $\mathbf{m}^*(\mathbf{p})$, which is an implicit function of $\mathbf{p}$ defined by $\mathbf{m}^* = T(\mathbf{m}^*; \mathbf{p})$. Solving the full bilevel program exactly is outside the scope of the proposed algorithm. Instead, we compute a first-order price direction and use multi-start empirical checks to assess the consistency of the attained stationary point. The finite-difference (FD) approach estimates the gradient by perturbing each price coordinate and re-solving the MFG, which is expressed as

$$\frac{\partial \Phi}{\partial p_k} \approx \frac{\Phi(\mathbf{p} + \epsilon \mathbf{e}_k) - \Phi(\mathbf{p})}{\epsilon}, \quad (39)$$

requiring $(K + 1) = 4$ MFG solutions per outer iteration. To address this computational bottleneck, we derive an exact analytical gradient using the IFT, reducing this to a single MFG solution with a lightweight 3×3 matrix computation.

1) *Per-User KKT Sensitivity:* To establish how the optimal allocation of each user responds to changes in the effective cost vector $\mathbf{c} = \mathbf{p} + \text{diag}(\boldsymbol{\eta}) \mathbf{m}$, we derive the following sensitivity result.

Lemma 1 (Per-User KKT Sensitivity). *For a user with channel $\mathbf{h}$, weight $\mathbf{w}$, and optimal allocation $\mathbf{b}^*$, let $\mathcal{A} \subseteq \{1, 2, 3\}$ denote the unconstrained active modalities ($b_k^* > b_k^{\min}$, i.e., neither inactive nor floor-binding). Then, the weighted curvature is defined as follows:*

$$D_k \triangleq w_k Q''_k(b_k^*, h_k) < 0, \quad k \in \mathcal{A}, \quad (40)$$

where the modality-specific second derivatives of the deployed capacity-compressed forms are given, with $z \triangleq (b \tilde{h})^p$, by

$$Q''_k(b, h) = -a_k \alpha_k (\alpha_k + 1) \kappa_k^2 \tilde{h}^2 \left(1 + \kappa_k b \tilde{h}\right)^{-(\alpha_k + 2)}, \quad k \in \{1, 2\}, \quad (41)$$

$$Q''_3(b, h) = a_3 \delta \tilde{h}^2 \text{sech}^2(\delta z) \left[p(p-1)(b \tilde{h})^{p-2} - 2\delta p^2 (b \tilde{h})^{2p-2} \tanh(\delta z) \right] \quad (42)$$

with $\tilde{h} = h\gamma$. The Jacobian $\mathbf{J} \in \mathbb{R}^{3 \times 3}$ with $J_{kj} = \partial b_k^* / \partial c_j$ satisfies the following conditions:

- (i) Budget inactive ($\lambda^* = 0$): $J_{kj} = \delta_{kj} / D_k$ for $k \in \mathcal{A}$ and zero otherwise.
- (ii) Budget active ($\sum_{k \in \mathcal{A}} b_k^* = B_0$): for $k, j \in \mathcal{A}$,

$$J_{kj} = \frac{\delta_{kj}}{D_k} - \frac{1}{D_k D_j S}, \quad S \triangleq \sum_{\ell \in \mathcal{A}} D_\ell^{-1}. \quad (43)$$

Proof. By differentiating the KKT condition in (22) with respect to c_j , we obtain

$$D_k \frac{\partial b_k^*}{\partial c_j} = \delta_{kj} + \frac{\partial \lambda^*}{\partial c_j}, \quad k \in \mathcal{A}. \quad (44)$$

In case (i), $\lambda^* = 0$ is constant; hence, $\partial \lambda^* / \partial c_j = 0$, and (44) directly yields the diagonal Jacobian. In case (ii), differentiating the binding budget constraint results in the following:

$$\sum_{k \in \mathcal{A}} \frac{\partial b_k^*}{\partial c_j} = 0. \quad (45)$$

By dividing (44) by D_k , summing over $k \in \mathcal{A}$, and applying (45), we have

$$\frac{\partial \lambda^*}{\partial c_j} = -\frac{1}{D_j S}, \quad (46)$$

from which (43) follows. For floor-binding modalities ($b_k^* = b_k^{\min}$), $J_{kj} = 0$ for all j , because the allocation is pinned. This completes the proof. $\square$

Remark 3 (Revenue-Adjusted Min-QoS Pricing). When the min-QoS floor is binding for modality k , the forced allocation $b_k^{\min}$ generates revenue $p_k b_k^{\min}$ regardless of price, creating a perverse incentive for the BS to inflate p_k . The revenue adjustment in (13) avoids this by crediting the BS only for the price-responsive portion of demand $(m_k^* - b_k^{\min})^+$. This adjustment propagates into the IFT gradient (49) via the revenue-adjusted mean-field vector $\mathbf{m}^{*,\text{adj}}$, ensuring the gradient does not reward price increases on floor-constrained modalities.

2) *IFT-Based Exact Gradient*: To derive the exact gradient, we define the fixed-point residual $\mathbf{F}(\mathbf{m}, \mathbf{p}) = \mathbf{m} - T(\mathbf{m}; \mathbf{p})$. As $\partial c_j / \partial p_l = \delta_{jl}$ and $\partial c_j / \partial m_l = \eta_j \delta_{jl}$, the Jacobians of the mean-field mapping T satisfy the following:

$$\frac{\partial T}{\partial \mathbf{p}} = \bar{\mathbf{J}}, \quad \frac{\partial T}{\partial \mathbf{m}} = \bar{\mathbf{J}} \text{diag}(\boldsymbol{\eta}), \quad (47)$$

where $\bar{\mathbf{J}} \triangleq \mathbb{E}_{\mathbf{h}, \mathbf{w}}[\mathbf{J}(\mathbf{h}, \mathbf{w})]$ denotes the population-averaged per-user Jacobian from Lemma 1.

Proposition 4 (IFT-Based Gradient). Let $\mathbf{A} \triangleq \mathbf{I} - \bar{\mathbf{J}} \text{diag}(\boldsymbol{\eta})$. Then, $\mathbf{A}$ is nonsingular (see Remark 4). The sensitivity of the equilibrium mean-field vector to prices is given by

$$\mathbf{M} \triangleq \frac{d\mathbf{m}^*}{d\mathbf{p}} = (\mathbf{I} - \bar{\mathbf{J}} \text{diag}(\boldsymbol{\eta}))^{-1} \bar{\mathbf{J}}, \quad (48)$$

Accordingly, the exact gradient of the BS objective (11) can be expressed as

$$\begin{aligned} \nabla_{\mathbf{p}} \Phi &= N \left[(1 - \rho)(\mathbf{m}^{*,\text{adj}} + \mathbf{M}^\top \mathbf{p}) \right. \\ &\quad \left. + \rho(\mathbf{I} + \mathbf{M}^\top \text{diag}(\boldsymbol{\eta})) \boldsymbol{\xi} \right], \end{aligned} \quad (49)$$

where $m_k^{*,\text{adj}} = (m_k^* - b_k^{\min})^+$ represents the revenue-adjusted mean-field vector from (13). Moreover, let $\boldsymbol{\xi}$ denote the welfare coupling vector, which is calculated as

$$\xi_l \triangleq \mathbb{E} \left[\sum_{k=1}^3 J_{kl} w_k Q'_k(b_k^*, h_k) \right] \quad (50)$$

capturing the sensitivity of the aggregate semantic quality to the cost perturbations. In the conventional operating regime, where $m_k^* > b_k^{\min}$ for all k , the adjustment decreases to $m_k^{*,\text{adj}} = m_k^* - b_k^{\min}$.

Proof. By applying the IFT to $\mathbf{F}(\mathbf{m}^*, \mathbf{p}) = \mathbf{0}$, we obtain

$$\left(\mathbf{I} - \frac{\partial T}{\partial \mathbf{m}} \right) \frac{d\mathbf{m}^*}{d\mathbf{p}} = \frac{\partial T}{\partial \mathbf{p}}, \quad (51)$$

yielding (48). The nonsingularity of $\mathbf{A}$ follows from its M -matrix structure (Remark 4).

To derive the gradient, we decompose the following:

$$\nabla_{\mathbf{p}} \Phi = \mathbf{g}_p + \mathbf{M}^\top \mathbf{g}_m, \quad (52)$$

where $\mathbf{g}_p = \partial \Phi / \partial \mathbf{p}|_{\mathbf{m}=\mathbf{m}^*}$ and $\mathbf{g}_m = \partial \Phi / \partial \mathbf{m}|_{\mathbf{p}}$ represent the partial gradients holding the other variable constant. For the revenue term $R = N \sum_k p_k (m_k^* - b_k^{\min})^+$ from (13), we have

$$g_{p,l}^R = N(m_l^* - b_l^{\min})^+, \quad g_{m,l}^R = N p_l \cdot \mathbb{I}_{\{m_l^* > b_l^{\min}\}}. \quad (53)$$

The indicator $\mathbb{I}_{\{m_l^* > b_l^{\min}\}}$ prevents the gradient from erroneously encouraging price increases on modalities with demand pinned at the QoS floor, where the marginal revenue with respect to m_l is zero. For the welfare term $W = N \mathbb{E}[\sum_k w_k Q_k(b_k^*)]$, because b_k^* depends on $\mathbf{p}$ and $\mathbf{m}$ through $\mathbf{c}$, we obtain

$$g_{p,l}^W = N \mathbb{E} \left[\sum_k w_k Q'_k J_{kl} \right] = N \xi_l, \quad g_{m,l}^W = N \eta_l \xi_l. \quad (54)$$

Combining the revenue and welfare terms yields

$$\begin{aligned} \mathbf{g}_p &= N[(1 - \rho)\mathbf{m}^{*,\text{adj}} + \rho \boldsymbol{\xi}], \\ \mathbf{g}_m &= N[(1 - \rho)\text{diag}(\boldsymbol{\sigma})\mathbf{p} + \rho \text{diag}(\boldsymbol{\eta})\boldsymbol{\xi}], \end{aligned}$$

where $\sigma_l = \mathbb{I}_{\{m_l^* > b_l^{\min}\}}$. In a typical operating regime, where $m_k^* > b_k^{\min}$ for all k , $\boldsymbol{\sigma} = \mathbf{1}$, and substituting it into $\mathbf{g}_p + \mathbf{M}^\top \mathbf{g}_m$ yields (49). This completes the proof. $\square$

When the budget is active, the welfare coupling vector simplifies to $\xi_l = \mathbb{E}[\sum_{k \in \mathcal{A}} J_{kl} c_k]$, eliminating the necessity of evaluating Q'_k in the outer loop. As a result, the IFT gradient requires only a single MFG solution per outer iteration (vs. 4 for FDs).

Remark 4 (Unconditional Well-Posedness of the IFT Matrix). The matrix $\mathbf{A} = \mathbf{I} - \bar{\mathbf{J}} \text{diag}(\boldsymbol{\eta})$ is nonsingular independently of the contraction condition (31). By Lemma 1, for budget-active users, the column sums of $\bar{\mathbf{J}}$ satisfy $\sum_{k \in \mathcal{A}} J_{kj} = 0$ (from (45)). For budget-inactive users, $\bar{\mathbf{J}}$ is diagonal with $J_{jj} = 1/D_j < 0$; hence, $\mathbf{1}^\top \bar{\mathbf{J}} \leq \mathbf{0}^\top$ componentwise, and $\mathbf{1}^\top \bar{\mathbf{J}} \leq \mathbf{0}^\top$. Consequently, the column sums of $\mathbf{A}$ satisfy $\sum_k A_{kj} = 1 - \eta_j \sum_k J_{kj} \geq 1 > 0$. As $D_k < 0$ and the

off-diagonal entries of $\mathbf{J}$ are nonnegative (from (43)), $\mathbf{A}$ has positive diagonal entries and nonpositive off-diagonal entries, making it a Z -matrix. A Z -matrix with strictly positive column sums is a nonsingular M -matrix, so $\mathbf{A}$ is invertible for all SNR regimes. Therefore, the contraction condition (31) guarantees MFG convergence [31], [32], whereas the IFT gradient is unconditionally well-posed.

Prices are updated using the projected Adam optimizer with the cosine-annealed learning rate $\alpha_{\text{lr}}^{(t)}$, with projections onto $[\underline{p}, \bar{p}]^3$. The moment averaging of the optimizer smooths residual MC noise in the Jacobian estimates. Convergence of the projected Adam optimizer on nonconvex bilevel objectives is not a certificate for the full bilevel program; accordingly, the returned price is interpreted as a stationary/KKT point supported by the IFT gradient and by multi-start empirical consistency.

D. Complete MFSE Algorithm

The complete MFSE algorithm (Algorithm 2) combines the inner MFG iteration with the outer Stackelberg pricing loop, incorporating the cosine-annealed Adam optimizer [33], Polyak price averaging [34] over the last W iterations, and an SNR-adaptive welfare weight. In particular, the welfare weight is given by

$$\rho(\gamma) = \rho_{\text{base}} + (1 - \rho_{\text{base}}) \cdot \exp(-\gamma/\gamma_0), \quad (55)$$

where $\gamma = 10^{\text{SNR}_{\text{dB}}/10}$ denotes the linear SNR, ρ_{base} represents a high-SNR welfare weight, and γ_0 controls the transition rate. The economic rationale is as follows: at a low SNR, where quality functions are in a steep region and where users derive a high marginal value from each unit of bandwidth, the operator prioritizes user welfare ($\rho \rightarrow 1$) to maintain service viability. At a high SNR, where diminishing returns occur and revenue extraction has a lower welfare cost, the operator reverts to a balanced revenue–welfare tradeoff ($\rho \rightarrow \rho_{\text{base}}$). This revenue–welfare leader objective follows the economic logic of two-layer Stackelberg resource allocation, in which an upper-layer provider balances its monetary return against the aggregate utility delivered to the lower layer [13]. Notably, the parameter ρ depends only on the system-level SNR (not on decision variables); hence, the IFT gradient (49) remains valid, with ρ treated as a fixed parameter per the SNR operating point. In all reported experiments we use $\rho_{\text{base}} = 0.85$ and $\gamma_0 = 10$; at the 20 dB operating SNR ($\gamma = 100$), (55) evaluates to $\rho(\gamma) = 0.850$, so the schedule reduces to the fixed value $\rho = 0.85$ reported in Table IV and departs from ρ_{base} only below ~ 13 dB (e.g., $\rho \approx 0.905$ at 10 dB). The schedule is therefore an optional low-SNR refinement, which is consistent with the ablation in Table VII that finds it inactive at the 20 dB test SNR.

E. Theoretical Analysis

In this subsection, we establish two critical theoretical properties of the proposed MFSE framework: the ε -Nash equilibrium guarantee and the computation and signaling costs.

Algorithm 2 MFSE with the Cosine-Annealed Adam Optimizer with Polyak Averaging

Require: Calibrated quality parameters $(\{a_k, \alpha_k, \kappa_k\}_{k=1,2}, a_3, \delta, p)$, congestion η , budget B_0 , channel scaling γ ; Adam optimizer parameters $(\alpha_{\text{max}}, \alpha_{\text{min}}, \beta_1, \beta_2)$; tolerances $(\varepsilon_{\text{mfg}}, \varepsilon_{\text{price}})$; Polyak window W ; MC samples $\{(\mathbf{h}^{(s)}, \mathbf{w}^{(s)})\}_{s=1}^{N_{\text{mc}}}$

Ensure: MFSE price $\mathbf{p}^*$ and mean-field vector $\mathbf{m}^*$

- 1: Initialize price $\mathbf{p}^{(0)}$; Adam optimizer states $\mathbf{v}^{(0)} = \mathbf{s}^{(0)} = \mathbf{0}$
- 2: Compute $\rho \leftarrow \rho(\gamma)$ via (55)
- 3: **for** $t = 0, 1, 2, \dots, T_{\text{outer}} - 1$ **do**
- 4: $\alpha_{\text{lr}}^{(t)} \leftarrow \alpha_{\text{min}} + \frac{1}{2}(\alpha_{\text{max}} - \alpha_{\text{min}})(1 + \cos(\pi t/T_{\text{outer}}))$
- 5: $\mathbf{m}^*, \{\mathbf{b}^{*(s)}\} \leftarrow \text{MFG-SOLVE}(\mathbf{p}^{(t)})$
- 6: **for** $s = 1, \dots, N_{\text{mc}}$ (**vectorized**) **do**
- 7: Compute $\mathbf{J}^{(s)}$ via Lemma 1
- 8: **end for**
- 9: $\bar{\mathbf{J}} \leftarrow \frac{1}{N_{\text{mc}}} \sum_s \mathbf{J}^{(s)}$; $\boldsymbol{\xi} \leftarrow \bar{\mathbf{J}}^\top \mathbf{c}$
- 10: $\mathbf{M} \leftarrow (\mathbf{I} - \bar{\mathbf{J}} \text{diag}(\eta))^{-1} \bar{\mathbf{J}}$
- 11: $\mathbf{g}^{(t)} \leftarrow (49)$ using $\mathbf{m}^*, \mathbf{M}, \boldsymbol{\xi}, \mathbf{p}^{(t)}, \rho$
- 12: Update $\mathbf{p}^{(t+1)}$ via the projected Adam optimizer with $\alpha_{\text{lr}} = \alpha_{\text{lr}}^{(t)}$
- 13: **end for**
- 14: $\mathbf{p}^* \leftarrow \frac{1}{W} \sum_{t=T_{\text{outer}}-W}^{T_{\text{outer}}-1} \mathbf{p}^{(t)}$
- 15: $\mathbf{m}^* \leftarrow \text{MFG-SOLVE}(\mathbf{p}^*)$
- 16: **return** $\mathbf{p}^*, \mathbf{m}^*$

1) ε -Nash Equilibrium Property: The principal theoretical contribution is that the MFSE, derived under the mean-field limit, offers an approximate equilibrium strategy profile for the finite N -player game.

Theorem 1 (ε -Nash Property of the MFSE). *Let $(\mathbf{p}^*, \mathbf{m}^*)$ be an MFSE, as in Definition 1, and let $\hat{\mathbf{b}}_i = \mathbf{b}^*(\mathbf{p}^*, \mathbf{m}^*, \mathbf{h}_i, \mathbf{w}_i)$ denote the allocation prescribed by the mean-field strategy for user i in the finite N -player game. Suppose the following conditions hold:*

- (A1) *The quality functions Q_k are uniformly bounded and Lipschitz on the active calibrated operating interval used by the finite game.*
- (A2) *The congestion coefficients satisfy $\eta_k > 0$ for all k .*
- (A3) *The channel gains and weights are independently and identically distributed (i.i.d.) across users with bounded fourth moments.*

Then, the strategy profile $(\hat{\mathbf{b}}_1, \dots, \hat{\mathbf{b}}_N)$ constitutes an ε -Nash equilibrium of the N -player Stackelberg game, that is, for every user i and every feasible deviation $\mathbf{b}'_i \in \mathcal{B}_i$,

$$U_i^{(N)}(\mathbf{b}'_i, \hat{\mathbf{b}}_{-i}) - U_i^{(N)}(\hat{\mathbf{b}}_i, \hat{\mathbf{b}}_{-i}) \leq \varepsilon(N), \quad (56)$$

where $U_i^{(N)}$ indicates the utility in the N -player game with an empirical mean-field vector, and

$$\varepsilon(N) = O\left(\frac{1}{\sqrt{N}}\right). \quad (57)$$

Proof. The proof comprises three steps.

Step 1 (Mean-field concentration). Under (A3), the allocations $\{\hat{\mathbf{b}}_j\}_{j=1}^N$ are i.i.d. with $\mathbb{E}[\hat{\mathbf{b}}_j] = \mathbf{m}^*$ by the consistency

of MFSE. Let the empirical mean-field distribution be defined as

$$\bar{m}_k^{(N)} \triangleq \frac{1}{N} \sum_{j=1}^N \hat{b}_{j,k}. \quad (58)$$

For each modality k , because $b_{j,k} \in [0, B_0]$, Hoeffding's inequality results in

$$\Pr\left(\left|\bar{m}_k^{(N)} - m_k^*\right| > t\right) \leq 2 \exp\left(-\frac{2Nt^2}{B_0^2}\right) \quad (59)$$

A union bound over $k = 1, 2, 3$ yields the following:

$$\left\|\bar{\mathbf{m}}^{(N)} - \mathbf{m}^*\right\|_\infty = O_p\left(1/\sqrt{N}\right). \quad (60)$$

Step 2 (Utility perturbation). For any user i with allocation $\mathbf{b}_i$, we define the empirical mean-field distribution when user i unilaterally deviates, whereas all others play $\hat{\mathbf{b}}_{-i}$:

$$\mathbf{m}^{(N)}(\mathbf{b}_i) \triangleq \bar{\mathbf{m}}^{(N)} + \frac{1}{N}(\mathbf{b}_i - \hat{\mathbf{b}}_i). \quad (61)$$

Because the utility difference is attributable solely from the congestion term,

$$\begin{aligned} & \left|U_i^{(N)}(\mathbf{b}_i, \hat{\mathbf{b}}_{-i}) - U_i(\mathbf{b}_i; \mathbf{p}^*, \mathbf{m}^*)\right| \\ & \leq \eta_{\max} B_0 \left\|\mathbf{m}^{(N)}(\mathbf{b}_i) - \mathbf{m}^*\right\|_\infty, \end{aligned} \quad (62)$$

where $\eta_{\max} = \max_k \eta_k$, this bound applies $\sum_k b_{i,k} \leq B_0$. Furthermore, $\mathbf{m}^{(N)}(\hat{\mathbf{b}}_i) = \bar{\mathbf{m}}^{(N)}$ (no shift when user i follows the prescribed strategy), whereas for a general deviation $\mathbf{b}'_i$,

$$\left\|\mathbf{m}^{(N)}(\mathbf{b}'_i) - \mathbf{m}^*\right\|_\infty \leq \left\|\bar{\mathbf{m}}^{(N)} - \mathbf{m}^*\right\|_\infty + \frac{B_0}{N}, \quad (63)$$

because $\|\mathbf{b}'_i - \hat{\mathbf{b}}_i\|_\infty \leq B_0$.

Step 3 (Deviation bound). For any deviation $\mathbf{b}'_i \in \mathcal{B}_i$, applying (62) to $\mathbf{b}'_i$ and $\hat{\mathbf{b}}_i$ results in the following:

$$\begin{aligned} & U_i^{(N)}(\mathbf{b}'_i, \hat{\mathbf{b}}_{-i}) - U_i^{(N)}(\hat{\mathbf{b}}_i, \hat{\mathbf{b}}_{-i}) \\ & \leq \underbrace{[U_i(\mathbf{b}'_i; \mathbf{p}^*, \mathbf{m}^*) - U_i(\hat{\mathbf{b}}_i; \mathbf{p}^*, \mathbf{m}^*)]}_{\leq 0 \text{ (by MF-optimality)}} \\ & \quad + \eta_{\max} B_0 \left\|\mathbf{m}^{(N)}(\mathbf{b}'_i) - \mathbf{m}^*\right\|_\infty \\ & \quad + \eta_{\max} B_0 \left\|\bar{\mathbf{m}}^{(N)} - \mathbf{m}^*\right\|_\infty \\ & \leq 2\eta_{\max} B_0 \left\|\bar{\mathbf{m}}^{(N)} - \mathbf{m}^*\right\|_\infty \\ & \quad + \frac{\eta_{\max} B_0^2}{N} = O(1/\sqrt{N}), \end{aligned} \quad (64)$$

where the final line substitutes (63) and incorporates the $O(1/N)$ mass-shift penalty into the dominant $O(1/\sqrt{N})$ concentration term from Step 1. This approach, which is standard for static MFGs [8], establishes (57). This completes the proof. $\square$

For the standard fixed-per-user-budget interpretation, B_0 , γ , and Q_k are independent of N , so the maximum per-user utility satisfies $U_{\max} = O(1)$ and the approximation ratio $\varepsilon/U_{\max} = O(1/\sqrt{N})$ vanishes as the population increases. Under the primary fixed-total model used in the revised evaluation, $B_0 = B_{\text{total}}/N$ and the same perturbation terms become smaller with N because each user's feasible mass

TABLE III: Computational Complexity Comparison

Method	Equil./train	Deploy.	Exchange
Finite Stack. [6]	$O(N^2 K)$	$O(NJ)$	model dep.
Central opt.	$O(NKJ)$	$O(NJ)$	full collect.
DRL	$O(NKT_{\text{tr}})$ offline	per-user $O(1)$	local/report
MFSE-IFT	$O(T_o T_m N_{\text{mc}} J)$	$O(NJ)$	9 scalars

shrinks. Theorem 1 therefore justifies the mean-field approach while keeping the bandwidth-scaling regime explicit. Although the MFG fixed point and user-level optimization each admit unique solutions, the outer Stackelberg objective (11) may have multiple local optima; the observed multi-start convergence is empirical evidence for a consistent stationary/KKT point, not a proof that the full bilevel maximum has been attained.

2) Complexity Analysis:

Theorem 2 (Computational Complexity). *The computational complexity of the MFSE algorithm (Algorithm 2) is as follows:*

- 1) **MFG solution (Algorithm 1):** Each fixed-point iteration requires $O(N_{\text{mc}} \cdot J)$ operations, where $J = O(\log(1/\varepsilon))$ represents the bisection work for the budget multiplier and the scalar speech inverse. The total MFG cost is denoted by $O(T_{\text{mfg}} \cdot N_{\text{mc}} \cdot J)$, which is independent of the number of users N for fixed N_{mc} .
- 2) **Stackelberg outer loop:** With the IFT gradients (Proposition 4), each outer iteration requires a single MFG solution with $O(N_{\text{mc}} K^2)$ Jacobian computations and a $K \times K$ matrix inversion. The total outer loop cost is denoted by $O(T_{\text{outer}} \cdot T_{\text{mfg}} \cdot N_{\text{mc}} \cdot J)$, a K -fold reduction over the FD approach (39), requiring $(K + 1)$ MFG solutions per iteration.
- 3) **Strategy deployment:** Applying the MFSE strategy to N users requires $O(N \cdot J)$ operations (one bisection per user).

As a result, the equilibrium computation is $O(1)$ in N for a fixed Monte Carlo sample size N_{mc} , whereas strategy deployment scales as $O(NJ)$. For signaling, the BS broadcasts the price vector $\mathbf{p} \in \mathbb{R}^3$ and mean-field vector $\mathbf{m} \in \mathbb{R}^3$, while each user uploads its demand $\mathbf{b}_i \in \mathbb{R}^3$. This requires a constant exchange of $2K + K = 9$ scalars per user per round, which is $O(1)$ in N . Compared to DRL, which also has $O(1)$ deployment complexity, MFSE provides explicit game-theoretic equilibrium, congestion-aware pricing, operator revenue optimization, and online re-solving without retraining.

Table III summarizes the resulting complexity and signaling comparison.

V. SIMULATION RESULTS

A. Simulation Setup

This section evaluates the proposed mean-field Stackelberg equilibrium (MFSE) under the deployed welfare-emphasis operating point $\rho = 0.85$, except for explicitly identified sensitivity studies. The primary scalability model fixes the total bandwidth at $B_{\text{total}} = 20$ MHz and assigns the per-user budget as $B_0(N) = B_{\text{total}}/N$. Thus, the default $N = 100$ setting corresponds to $B_0 = 0.2$ MHz. Table IV summarizes the main simulation parameters.

TABLE IV: Simulation Parameters

Parameter	Value	Parameter	Value
Users N	100	Modalities K	3
Total BW B_{total}	20 MHz	B_0 at $N = 100$	0.2 MHz
SNR	20 dB	N_{mc}	2000
Damping θ	0.3	MFG tol.	10^{-6}
Welfare ρ_{base}	0.85	Price bounds	$[0.01, 5.0]$
γ_0 (adaptive ρ)	10	Adaptive ρ	optional
Channel	Rayleigh	Weights $\mathbf{w}$	$(1, 1, 1) \pm 0.3$
η	$(0.10, 0.15, 0.12)$	$\mathbf{b}^{\min}$	$(0, 0.01, 0)$ MHz
Image	Top-1 accuracy	Text/speech	BLEU/PESQ

The MFSE is compared with the following seven baselines encompassing centralized, distributed, DRL, and game-theoretic alternatives:

- *Centralized semantic optimization (Sem-Opt)*: This is a pre-sliced centralized semantic optimizer that serves as the welfare upper bound under the per-user bandwidth partition.
- *Equal allocation*: This alternative distributes the bandwidth uniformly ($b_{i,k} = B_0/K$) to every modality.
- *PPO, DDPG, and SAC*: These three DRL baselines train per-user multilayer perceptron (MLP) policies mapping the local six-dimensional state $[w_i, h_i]$ to softmax-constrained allocations with Dirichlet exploration, applying early stopping (a patience of 50 episodes and a threshold of 0.1%) for up to 1000 episodes, where PPO is retrained on the capacity-compressed reward and DDPG/SAC are auxiliary references.
- *Distributed alternating direction method of multipliers (ADMM)*: This distributed baseline uses the same price, mean-field, and per-user demand messages as MFSE but updates prices using dual sub-gradient tracking.
- *Finite Stackelberg*: This comparator solves the exact N -player game with multistart initialization and analytic pricing gradients, using Stackelberg and implicit-differentiation tools consistent with [6], [31], [32].

B. Calibration and Metrics

To calibrate the capacity-compressed quality functions of the multi-modal semantic communication system, we utilize real codec measurements. Specifically, the image quality is calibrated using the downstream CIFAR-10 classifier top-1 accuracy from Deep JSCC reconstructions. The text and speech qualities are calibrated using the DeepSC BLEU score over the Europarl dataset and the DeepSC-S PESQ score over the VoiceBank-DEMAND dataset, respectively. Without loss of generality, we adopt the following parametric models for the quality functions:

$$Q_k(b, \text{SNR}) = a_{\max} [1 - (1 + \kappa b \text{SNR})^{-\alpha}], \quad k \in \{1, 2\}, \quad (65)$$

and

$$Q_{\text{speech}}(b, \text{SNR}) = a_{\max} \tanh(\delta(b \text{SNR})^p), \quad (66)$$

where b denotes the allocated bandwidth. The fitted parameters are determined as $(a_{\max}, \alpha, \kappa) = (0.8991, 1.3819, 8.0221)$ for image, $(0.3403, 4.2186, 0.7017)$ for text, and $(a_{\max}, \delta, p) = (0.3136, 1.0901, 0.3569)$ for speech. These models yield high




Fig. 2: Real-codec calibration curves.

goodness of fit, with reference-rate coefficients of determination $R^2 = 0.9987$ (image), 0.9944 (text), and 0.9999 (speech), while every per-rate calibration point satisfies $R^2 \geq 0.994$. Accordingly, the measured and fitted calibration curves are illustrated in Fig. 2.

To quantify the overall semantic performance, we define the smooth weighted semantic quality (WSQ) metric as

$$\bar{Q} = \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K w_{i,k} Q_k(b_{i,k}, h_{i,k}), \quad (67)$$

where $w_{i,k}$ is user i 's preference weight for modality k . To evaluate the literal end-to-end performance, the allocated bandwidth is snapped to the nearest trained codec rate. The normalized end-to-end score is computed as the average of the image accuracy, BLEU/100, and $(\text{PESQ} - 1)/3.5$. Furthermore, we evaluate the fairness of the resource allocation using Jain's fairness index, which is calculated as

$$\mathcal{J} = \frac{\left(\sum_{i=1}^N q_i\right)^2}{N \sum_{i=1}^N q_i^2}, \quad (68)$$

where $q_i = \sum_{k=1}^K w_{i,k} Q_k(b_{i,k}, h_{i,k})$ denotes the WSQ of user i .

To evaluate task-priority steering (Remark 2), we assess downstream task performance under welfare weight $\rho = 0.90$ and priority vector $\pi = (1, 2, 1)$, steering resources to the image modality. As shown in Fig. 3, MFSE achieves the




Fig. 3: End-to-end literal task metrics by modality.

TABLE V: Smooth WSQ comparison across baselines.

Method	WSQ (% of Sem-Opt)	95% CI
MFSE (ours)	98.89	[98.81, 98.98]
Finite Stackelberg	98.92	[98.86, 98.99]
PPO	98.80	[98.70, 98.90]
ADMM	98.69	[98.62, 98.76]
SAC	98.06	[97.89, 98.22]
Equal	97.87	[97.77, 97.98]
DDPG	96.88	[96.50, 97.27]

highest normalized end-to-end score of 0.3519, outperforming PPO (0.3422), Sem-Opt (0.3265), ADMM (0.3104), and equal allocation (0.2769). This corresponds to improvements of 2.82%, 13.36%, and 27.10% over PPO, ADMM, and equal allocation, respectively. Notably, MFSE dominates PPO across all modalities: image accuracy (0.660 vs. 0.646), text BLEU (13.59 vs. 12.56), and speech PESQ (1.909 vs. 1.892), demonstrating effective priority steering with $O(1)$ complexity.

Table V reports the smooth WSQ across baselines over eight seeds, with paired per-seed differences used to assess statistical significance. The proposed MFSE attains 98.89% of the centralized Sem-Opt welfare and is statistically tied with the deployable fair-DRL (PPO) policy (98.80%): the paired MFSE–fair-DRL difference has a 95% CI $[-0.000008, 0.002748]$ that crosses zero, so the two match in welfare rather than one dominating. MFSE also matches the exact finite-player Stackelberg game (98.92%) within statistical noise; Section V-C quantifies a maximum WSQ gap of only 0.19%, confirming the accuracy of the mean-field approximation while avoiding its computational cost. By contrast, MFSE significantly exceeds ADMM (98.69%), SAC (98.06%), and equal allocation (97.87%), whose paired MFSE–X differences are all strictly positive (for example, +0.00301 and +0.01535 for ADMM and equal allocation, respectively). Hence, MFSE matches the deployed DRL agent’s welfare while additionally providing an ε -Nash guarantee, congestion-aware pricing, and $O(1)$ per-user cost.

We sweep the welfare weight ρ to analyze the operator revenue versus user semantic quality tradeoff. As shown in Fig. 4, the default weight allows MFSE to balance goals, preserving a positive BS revenue of 3.647 while attaining 98.99% of the centralized WSQ. Hence, we adopt this value as the compromise operating point.




Fig. 4: Welfare-revenue operating-point sweep.


(a) Constant-LR Adam ($\alpha_{lr} = 0.1$). (b) Cosine-annealed Adam with Polyak averaging.

Fig. 5: Stackelberg price trajectories versus outer iterations.

C. Convergence and Gradient Efficiency

To evaluate the convergence behavior of the Stackelberg price optimization, we compare a constant-learning-rate Adam optimizer against a cosine-annealed Adam optimizer combined with Polyak averaging [33], [34] in Fig. 5. Under the default welfare operating point ρ , the constant-learning-rate updates settle near the oscillating price vector $\mathbf{p} = (0.36, 0.27, 0.23)$. By contrast, the cosine-annealed scheme converges smoothly in 15 outer iterations to the stable equilibrium price vector $\mathbf{p}^* = (0.27, 0.19, 0.15)$. Notably, the analytical implicit function theorem (IFT) gradient remains highly consistent with the finite-difference check, thereby avoiding the necessity of repeated mean-field game (MFG) solves.

To analyze the mathematical validity of the fixed-point iteration, we evaluate the contraction condition (31) under different channel conditions. In the high SNR regime (15–30 dB), the contraction condition holds comfortably, with the row-sum norm $\|\bar{\mathbf{J}} \text{diag}(\boldsymbol{\eta})\|_{\infty}$ lying within the range of 0.004–0.006. Conversely, at a lower SNR of 10 dB, the undamped sufficient condition is not satisfied, yielding a value of 4.39. Fortunately, the damped update scheme in (36) with damping factor $\theta = 0.3$ yields a spectral radius of approximately 0.70, thereby explaining the robust convergence observed in the low-curvature regime.

To validate the accuracy of the proposed mean-field game approximation, Fig. 6a compares the smooth WSQ of MFSE against the best multi-start finite-population Stackelberg game. Across user populations $N \in \{5, 10, 20, 50\}$, the maximum absolute WSQ gap between MFSE and the best multi-start finite-Stackelberg solution is only 0.1864% under the default welfare weight ρ . This confirms that MFSE matches the exact finite-population game closely while Sem-Opt provides the centralized upper bound. However, the two approaches exhibit a substantial difference in computational cost.




Fig. 6: System welfare comparison under finite-player game scaling and cross-modal correlation.




Fig. 7: Computation time versus user population.

To demonstrate the scalability of the proposed framework, Fig. 7 presents the wall-clock computation time of all schemes as the user population N increases. In particular, the solve time of MFSE remains the lowest and nearly flat at approximately 0.2–0.3 s across all N , which is consistent with the theoretical $O(1)$ per-user equilibrium complexity. In contrast, the exact finite-Stackelberg solver scales the most steeply and leaves the plotted range at moderate populations, since it must repeatedly solve the full N -player game with multiple restarts. The remaining baselines lie in between: the pre-sliced Sem-Opt and the DDPG policy grow with N to roughly one second, whereas the PPO and SAC policies incur higher and more fluctuating inference costs of up to a few seconds; the equal-allocation heuristic is trivially inexpensive but solves no game. Consequently, MFSE attains the same semantic quality as the exact finite-population model (Fig. 6a) while sustaining the lowest and most stable computational cost among the compared schemes.

D. Priority and Cross-Modal Sensitivity

To evaluate the sensitivity of the resource allocation to semantic priorities, we sweep the image priority multiplier $\pi_{i,\text{image}}$ from 0.5 to 5.0 under the test welfare weight $\rho = 0.7$, while keeping the text and speech priorities fixed at 1.0. As shown in Fig. 8, increasing the image priority multiplier causes the allocated image bandwidth to rise from 0.0397 to 0.0566 MHz, which corresponds to a 42.64% increase. Accordingly, the image bandwidth share increases from 0.260 to 0.372. Notably, the original-weight WSQ remains nearly flat and non-monotone during this sweep. Therefore, this priority experiment demonstrates that the proposed framework




Fig. 8: Bandwidth allocation and prices under priority steering.

TABLE VI: Fixed-total MFSE scalability.

N	10	20	50	100	200	500	1000
B_0 (MHz)	2.00	1.00	0.40	0.20	0.10	0.04	0.02
WSQ	1.452 ± 0.090	1.520 ± 0.082	1.533 ± 0.033	1.491 ± 0.013	1.420 ± 0.020	1.297 ± 0.004	1.179 ± 0.015

achieves controllable bandwidth steering while maintaining overall system quality.

To investigate the impact of cross-modal redundancy on system performance, we evaluate the correlation model extension using the effective bandwidth relation $b^{\text{eff}} = R(\rho_c)b$ under the test welfare weight $\rho = 0.7$, where ρ_c represents the cross-modal correlation strength. As illustrated in Fig. 6b, at a correlation strength of $\rho_c = 0.8$, the proposed scheme achieves an effective WSQ gain of 1.49% and an effective-bandwidth saving of 11.84% compared to the independent coding case. Consequently, this bandwidth saving reflects the system’s ability to exploit shared semantic content to avoid redundant transmission across modalities.

E. Unified Performance and Scalability Comparison

To evaluate the system performance under varying channel conditions, we plot the WSQ and Jain’s fairness index across a wide range of SNR values (0–30 dB) in Fig. 9a and Fig. 9b, respectively. At the default SNR of 20 dB, MFSE achieves a WSQ of 1.4912. Under the same configuration, the Jain fairness index values are 0.9588 for MFSE, 0.9585 for Sem-Opt, and 0.9565 for equal allocation, as shown in Fig. 9b. Consequently, the fairness performance is highly comparable across all optimized resource allocation schemes. Furthermore, under the legacy constant-per-user regime (retained only to validate the clean $O(1/\sqrt{N})$ rate of Theorem 1), Fig. 10 illustrates the ε -Nash gap and the computational time versus the user population N . Notably, the structural ε -Nash gap decreases monotonically from 2.94×10^{-7} at $N = 10$ to 1.55×10^{-8} at $N = 1000$, which is consistent with the theoretical $O(1/\sqrt{N})$ rate established in Theorem 1. Additionally, the IFT solver achieves a constant execution time of approximately 0.25 s, whereas the finite-difference approach requires 1.27 s.

To evaluate the system scalability under resource constraints, we investigate a fixed-total spectrum scenario where the total bandwidth is set to $B_{\text{total}} = 20$ MHz and the per-user budget is scaled as $B_0 = B_{\text{total}}/N$. The resulting WSQ values are summarized in Table VI. Under this scaling, the per-user budget shrinks as the user population grows, which leads to resource dilution. For user populations $N \leq 100$, the measured fixed-total ε -Nash gaps are 3.54×10^{-4} , 9.32×10^{-5} ,





Fig. 9: System performance metrics versus operating SNR (0–30 dB).



(a) ϵ -Nash gap vs. N with theoretical $O(1/\sqrt{N})$ reference. (b) MFSE computation time: IFT (~ 0.25 s) vs. FD (~ 1.27 s), both $O(1)$ in N .

Fig. 10: Legacy constant-per-user ablation, retained to validate the $O(1/\sqrt{N})$ rate of Theorem 1.

3.93×10^{-6} , and 6.14×10^{-8} for $N = 10, 20, 50$, and 100 , respectively, which validates the high accuracy of the mean-field approximation even under limited spectrum resources.

Fig. 11 compares MFSE against centralized and DRL baselines under the fixed-total model. Using a common Monte Carlo population, MFSE approaches the centralized Sem-Opt upper bound and matches or exceeds all decentralized and learned baselines for $N \geq 50$, with the margin widening as the population grows (e.g., exceeding PPO by $+0.13\%$ at $N = 50$ and $+0.55\%$ at $N = 1000$, and equal allocation by $+3.5\%$ at $N = 1000$). Congestion-internalizing pricing guides the decentralized users to the social optimum, whereas congestion-blind baselines degrade under resource scarcity. At $N = 10$, MFSE trails Sem-Opt by only 0.56% due to stronger finite-player coupling. Both MFSE and DRL have $O(1)$ deployment complexity, but MFSE provides an ϵ -Nash equilibrium, congestion-aware pricing, revenue maximization, and online adaptation with a small overhead of nine scalars per round, avoiding offline training.

For $N \geq 200$, equilibrium prices approach the upper bound ($p_{\max} = 5.0$) as the per-user budget $B_0 = B_{\text{total}}/N$ shrinks, which reflects increased congestion. Even at this boundary, the ϵ -Nash gap remains negligible, verifying the mean-field approximation. For larger populations, price bounds can be widened to ensure interior convergence.

F. Ablation Study

To evaluate individual components, Table VII presents an ablation study under a tight budget of $T_{\max} = 24$ outer iterations at 20 dB SNR. Although all variants eventually




Fig. 11: Weighted semantic quality versus user population.

TABLE VII: Ablation study of MFSE components.

Variant	WSQ	Fairness	Conv. (%)	Time (s)
Full MFSE	1.481 ± 0.007	0.9586	100	5.2
No cosine annealing	1.481 ± 0.007	0.9586	29	7.0
No Polyak averaging	1.480 ± 0.007	0.9585	100	4.9
No adaptive ρ	1.481 ± 0.007	0.9586	100	5.2
No min-QoS floors	1.479 ± 0.007	0.9585	100	5.3
No IFT (finite diff.)	1.478 ± 0.007	0.9585	100	19.7

converge to the same equilibrium, the tight budget highlights the convergence rate and solve time differences. The full MFSE achieves 100% convergence while securing top-tier WSQ and fairness.

Removing cosine annealing drops convergence to 29% due to oscillations (Fig. 5). Disabling Polyak averaging or QoS floors reduces WSQ and fairness. Replacing the analytical IFT gradient with finite-differences increases solve time from 5.2 s to 19.7 s. The adaptive- ρ schedule is inactive at this high test SNR. Overall, full MFSE is the only configuration achieving high quality, fairness, and convergence within tight iterations.

VI. CONCLUSION

In this work, we investigated the resource allocation problem in scalable multimodal SemCom networks. Specifically, we proposed a mean-field Stackelberg game framework where the BS acts as the leader to set modality-specific prices, while the users optimize their bandwidth demands against the mean-field aggregate behavior. The proposed game-theoretic approach addresses the coupling of user decisions, making the framework feasible for large-scale deployment. For future research, we plan to extend this framework by considering multicell coordination with intercell interference, dynamic user arrivals and departures, and joint source–channel coding under semantic-aware quality constraints.

REFERENCES

- [1] C. Chaccour, W. Saad, M. Debbah, Z. Han, and H. Vincent Poor, “Less data, more knowledge: Building next-generation semantic communication networks,” *IEEE Communications Surveys & Tutorials*, vol. 27, no. 1, pp. 37–76, 2025.
- [2] W. Xie, T. Zhang, M. Xiong, J. Zou, P. Zhu, and L. Yang, “Multimodal semantic communication: Research progress, key challenges, and future trends,” *IEEE Communications Standards Magazine*, vol. 9, 12 2025. [Online]. Available: <http://dx.doi.org/10.1109/mcomstd.2025.3579425>
- [3] M. Chen, P. Yan, H. Li, C. Li, S. Dang, and A. Nallanathan, “Task-oriented resource allocation for cross-modal semantic communications,” *IEEE Network*, pp. 1–8, 2025.

- [4] L. Yan, Z. Qin, R. Zhang, Y. Li, and G. Y. Li, "Resource allocation for text semantic communications," *IEEE Wireless Communications Letters*, vol. 11, no. 7, pp. 1394–1398, 2022.
- [5] B. Guler, A. Yener, and A. Swami, "The semantic communication game," *IEEE Transactions on Cognitive Communications and Networking*, vol. 4, 12 2018. [Online]. Available: <http://dx.doi.org/10.1109/tccn.2018.2872596>
- [6] J. Liu, M. Xiao, J. Wen, J. Kang, R. Zhang, T. Zhang, D. Niyato, W. Zhang, and Y. Liu, "Optimizing resource allocation for multi-modal semantic communication in mobile aigc networks: A diffusion-based game approach," *IEEE Transactions on Cognitive Communications and Networking*, vol. 11, no. 5, pp. 3346–3360, 2025.
- [7] S. Aggarwal, M. A. u. Zaman, M. Bastopcu, and T. Başar, "Semantic communication in multi-team dynamic games: A mean field perspective," *IEEE Transactions on Automatic Control*, vol. 71, no. 1, pp. 49–64, 2026.
- [8] M. Huang, R. P. Malhame, and P. E. Caines, "Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle," *Communications in Information and Systems*, vol. 6, no. 3, pp. 221–252, 2006.
- [9] J.-M. Lasry and P.-L. Lions, "Mean field games," *Japanese Journal of Mathematics*, vol. 2, no. 1, pp. 229–260, 2007.
- [10] Y. Xiao, X. Zhang, Y. Li, G. Shi, and T. Başar, "Rate-distortion theory for strategic semantic communication," in *2022 IEEE Information Theory Workshop (ITW)*, 2022, pp. 279–284.
- [11] C. K. Thomas and W. Saad, "Hypergame theory for decentralized resource allocation in multi-user semantic communications," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 2024, pp. 6036–6043.
- [12] X. Tang, Z. Ma, B. Li, C. Li, Q. Du, D. Niyato, and Z. Han, "Active ris-aided anti-jamming wireless communications: A stackelberg game perspective," *IEEE Transactions on Communications*, vol. 74, pp. 2612–2625, 2026.
- [13] T. Lyu, H. Xu, F. Liu, M. Li, L. Li, and Z. Han, "Two layer stackelberg game-based resource allocation in cloud-network convergence service computing," *IEEE Transactions on Cognitive Communications and Networking*, vol. 10, no. 6, pp. 2412–2426, 2024.
- [14] M. Tiken Singh, A. Arif, R. Kumar Prasad, B. Choudhury, C. Kalita, and S. M. S. Askari, "Optimizing fairness and spectral efficiency with shapley-based user prioritization in semantic communication," *IEEE Access*, vol. 13, pp. 61 299–61 311, 2025.
- [15] R. Fantacci and B. Picano, "Multi-user semantic communications system with spectrum scarcity," *Journal of Communications and Information Networks*, vol. 7, no. 4, pp. 375–382, 2022.
- [16] X. Yang, H. Yang, Y. Jiang, A. Alphones, and L. Xiao, "Game-guided matching theory-based resource allocation for secure semantic communications," *IEEE Transactions on Vehicular Technology*, vol. 74, no. 5, pp. 8357–8362, 2025.
- [17] Z. Ji, D. Wu, X. Shen, and X. Guan, "Joint extraction ratio, mode selection, and channel allocation for semantic-based d2d communications," *IEEE Transactions on Cognitive Communications and Networking*, vol. 12, pp. 3913–3927, 2026.
- [18] R. Zeng, C. Li, L. Li, M. Li, and T. Hu, "Solving mean field game based on physics-informed operator learning and deep reinforcement learning," *Journal of Computational Physics*, vol. 545, p. 114457, Jan. 2026.
- [19] Z. Su, Q. Qi, Q. Xu, S. Guo, and X. Wang, "Resource allocation in wireless powered IoT system: A mean field Stackelberg game-based approach," *Sensors*, vol. 18, no. 10, p. 3173, 2018.
- [20] C. Wang and B. L. Mark, "Hierarchical mean field/Stackelberg game power control for D2D cooperative relaying networks," in *IEEE Global Communications Conference (GLOBECOM)*, 2022, pp. 2855–2860.
- [21] A. Benamor, O. Habachi, I. Kammoun, and J.-P. Cancès, "Mean field game-theoretic framework for distributed power control in hybrid NOMA," *IEEE Transactions on Wireless Communications*, vol. 21, no. 12, pp. 10 920–10 935, 2022.
- [22] L. Li, Q. Cheng, X. Tang, T. Bai, W. Chen, Z. Ding, and Z. Han, "Resource allocation for noma-mec systems in ultra-dense networks: A learning aided mean-field game approach," *IEEE Transactions on Wireless Communications*, vol. 20, no. 3, pp. 1487–1500, 2021.
- [23] H. Zhang, C. Lu, H. Tang, X. Wei, L. Liang, L. Cheng, W. Ding, and Z. Han, "Mean-field-aided multiagent reinforcement learning for resource allocation in vehicular networks," *IEEE Internet of Things Journal*, vol. 10, no. 4, pp. 3490–3505, 2023.
- [24] S.-J. Lee and Y. H. Lee, "Semantic communications - a comprehensive survey for future research issues," in *2024 Fifteenth International Conference on Ubiquitous and Future Networks (ICUFN)*, 2024, pp. 660–664.
- [25] I. Ahmed, Y. Sun, J. Fu, A. Köse, L. Musavian, M. Xiao, and B. Özbek, "Semantic communications in 6g: Coexistence, multiple access, and satellite networks," *IEEE Communications Standards Magazine*, vol. 9, 12 2025. [Online]. Available: <http://dx.doi.org/10.1109/mcomstd.2025.3592327>
- [26] C. Zhang, L. Huang, and Q. Ning, "Resource allocation in wireless semantic communications: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, 2026.
- [27] H. Hu, X. Zhu, F. Zhou, W. Wu, R. Qingyang Hu, and H. Zhu, "Resource allocation for multi-modal semantic communication networks," *IEEE Transactions on Communications*, vol. 73, no. 9, pp. 7599–7616, 2025.
- [28] L. Xia, Y. Sun, D. Niyato, L. Zhang, and M. Ali Imran, "Wireless resource optimization in hybrid semantic/bit communication networks," *IEEE Transactions on Communications*, vol. 73, no. 5, pp. 3318–3332, 2025.
- [29] C. Xu, M. B. Mashhadi, Y. Ma, R. Tafazolli, and J. Wang, "Generative semantic communications with foundation models: Perception-error analysis and semantic-aware power allocation," *IEEE Journal on Selected Areas in Communications*, vol. 43, no. 7, pp. 2493–2505, 2025.
- [30] S. Hua, Y. Sun, K. Ma, L. Feng, M. Chen, Z. Yang, and M. A. Imran, "Bandwidth management in semantic communications: A tradeoff between data sensing and transmission," *IEEE Transactions on Vehicular Technology*, pp. 1–13, 2025.
- [31] T. Fiez, B. Chasnov, and L. J. Ratliff, "Implicit learning dynamics in stackelberg games: Equilibria characterization, convergence analysis, and empirical study," in *Proceedings of the 37th International Conference on Machine Learning (ICML)*. PMLR, 2020, pp. 3133–3143.
- [32] J. Lorraine, P. Vicol, and D. Duvenaud, "Optimizing millions of hyperparameters by implicit differentiation," in *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, 2020, pp. 1540–1552.
- [33] I. Loshchilov and F. Hutter, "SGDR: Stochastic gradient descent with warm restarts," in *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.
- [34] B. T. Polyak and A. B. Juditsky, "Acceleration of stochastic approximation by averaging," *SIAM Journal on Control and Optimization*, vol. 30, no. 4, pp. 838–855, 1992.