

Chapter 21

Intermediation in Online Advertising



Santiago R. Balseiro, Ozan Candogan, and Huseyin Gurkan

Abstract In online advertising, impressions are sold via real-time auctions which are organized by central platforms referred to as ad exchanges. For technological or operational reasons, advertisers generally participate in the auctions run by exchanges through intermediaries which acquire impressions on their behalf. Intermediaries are specialized entities that provide targeted services for a particular segment of the market, and typically there are multiple stages of intermediation. Moreover, an advertiser may have private information, e.g., budget, targeting criterion or value attributed to an impression. The presence of intermediaries and this information asymmetry introduce several new research questions. In the first part of this chapter, we study the mechanism design problem of an intermediary who offers a contract to an advertiser with a private budget and a private targeting criterion. We characterize the optimal mechanism and establish that the presence of the intermediary results in simpler bidding policies. In the second part of this chapter, we study the strategic interaction among intermediaries organized in a chain network. We characterize a subgame perfect equilibrium of the resulting game among intermediaries and show that the most profitable position in the intermediation chain depends on the underlying value distribution of the advertiser.

S. R. Balseiro (✉)
Columbia University, New York, NY, USA
e-mail: srb2155@columbia.edu

O. Candogan
University of Chicago, Chicago, IL, USA
e-mail: ozan.candogan@chicagobooth.edu

H. Gurkan
Duke University, Durham, NC, USA
e-mail: hg67@duke.edu

21.1 Introduction

Online advertising is a rapidly growing market whose annual revenue exceeded 72.5 billion dollars in the United States in 2016 (Internet Advertising Bureau 2016). This growth has been accompanied by technological advancements which introduce novel tactical and operational challenges for both publishers (the supply side) and advertisers (the demand side) such as real-time bidding and sophisticated targeting. To overcome these challenges, publishers and advertisers increasingly work with intermediaries who have emerged to facilitate transactions between two sides by providing technological and managerial services.

Specifically, when a user visits a publisher's page, an advertising opportunity, referred to as an impression, is generated. This impression is supplied to an exchange, where they are auctioned. These auctions take place in milliseconds after the user's visit to the web page. Due to the real-time nature of these auctions, participants generally employ sophisticated algorithms for automatically targeting users based on specific metrics. Many advertisers work with advertising agencies, who often focus on serving similar clients, enabling them to become experts at campaign management for a specific industry (e.g., pharmaceutical, automotive). Advertising agencies are service-based organizations, and they serve as a managerial layer, typically on top of a licensed demand side platform. Demand side platforms (DSPs) aggregate demand from different market participants and provide real-time bidding service in the auctions run by the exchanges. In addition, there are other intermediaries who specialize in services such as re-targeting (tracking a particular impression in different websites), measurement, and analytics. These intermediaries support the entire trading infrastructure.

The presence of intermediaries in this industry introduces several new interesting questions. What kind of contracts should an intermediary offer to an advertiser? How should an intermediary bid on behalf of its customers? How does the presence of an intermediary affect the efficiency of the market? How does the structure of the intermediation network affect the profits of its participants? Do intermediaries prefer to be closer to the supply source or demand source? This chapter sheds light on these issues by reviewing two separate models which are studied by Balseiro and Candogan (2017) and Balseiro et al. (2017). In the first model (hereafter OCI), we characterize the optimal contract offered by an intermediary to an advertiser in a setting where the advertiser's budget and targeting criteria are private. Using our results, we show that the presence of the intermediary results in simpler bidding policies. In the second model (hereafter MSI), we focus on a multi-stage intermediation network, and analyze the relation between the intermediation network and the profits of the market participants. To do so, we provide a game theoretic model where intermediaries in a chain network between an exchange and an advertiser with private values, sequentially select their mechanisms from a practically relevant class of mechanisms. We characterize a subgame perfect equilibrium for this game, and show that the most profitable position in the intermediation chain depends on the underlying value distribution of the advertiser. In the remainder of this section, we detail our contributions and relate them to the existing literature.

21.1.1 Main Contributions

In the OCI model (see Sect. 21.2), we study the dynamic mechanism design problem of an intermediary who offers a contract to an advertiser with a private budget and targeting criteria. Since the private information of the advertiser is multi dimensional, this problem in general is hard to solve. Therefore, we develop a novel solution method which combines a performance space characterization technique and a duality-based approach. Specifically, we first characterize the performance space that consists of the expected cost and value achievable by any feasible (dynamic) bidding policy, and use our duality-based approach to reduce the optimal contract design problem to a tractable convex optimization problem. In this way, we obtain a crisp characterization of the intermediary's optimal bidding policy. The policy is stationary and has two notable features: (i) the policy bids a weighted average of the values associated with different types, and (ii) the bids are appropriately shaded. Here, bidding a weighted average of the values ensures truthful reporting of the advertiser while bid shading accounts for budget constraints. Using our results, we establish that the intermediary can profitably provide bidding service to a budget-constrained advertiser, and in some cases increase the overall market efficiency.

Differently from the OCI model, in the MSI model (see Sect. 21.3), we consider a setting where an advertiser seeks to acquire an impression from an exchange through a chain of intermediaries. Using a game theoretic model, we study the mechanisms offered by intermediaries when the advertiser's value is private. We characterize a subgame perfect equilibrium of the game between intermediaries within the class of second-price mechanisms which are commonly used in the display advertising market. We show that economic incentives are not necessarily aligned along the chain, i.e., profit-maximizing intermediaries have incentives to shade bids and not to allocate impressions, even when profitable for their downstream customers. Moreover, we establish that the position in the intermediation network has a significant impact on the profits of the intermediaries, and the most profitable position depends on the underlying value distribution of the advertiser.

The proofs of the claims in Sects. 21.2 and 21.3 can be found in Balseiro and Candogan (2017) and Balseiro et al. (2017), respectively.

21.1.2 Literature Review

The models considered in this chapter contribute to various streams of literature, namely to those of intermediary problems, online advertising, and mechanism design with budget constraints.

21.1.2.1 Intermediary Problems

Feldman et al. (2010) and Stavrogiannis et al. (2013) focus on settings where captive buyers bid through intermediaries to acquire impressions. In these papers, the value distribution of the buyer has bounded support, and intermediaries directly bid at the exchange. As opposed to the OCI model, these papers focus on captive advertisers who are not liquidity constrained. Moreover, in these studies there are no “multiple intermediation tiers” whereas the MSI model analyzes how multiple intermediaries share surplus under different value distributions. Additionally, in these papers, intermediaries are restricted to forwarding the highest bid they receive from the buyers. However, in our models, we show that at equilibrium strategic intermediaries shade their bids, as opposed to simply reporting upstream the highest downstream bid. Loertscher and Niedermayer (2007, 2012) and Niazadeh et al. (2014) consider fee setting mechanisms for an intermediary with the two-sided private information setting introduced by Myerson and Satterthwaite (1983). These papers consider a single intermediary (as opposed to the MSI model) with a captive seller and captive buyer without budget constraints (as opposed to the OCI model), and provide conditions under which an affine fee structure that is commonly used in practice is optimal. In another recent work, Manea (2018) focuses on a setting where intermediaries trade a single good in a network, by considering a complete information setting in which buyers’ values are deterministic and common knowledge in contrast to our models. We refer the reader to Condorelli and Galeotti (2016) for a review of the recent literature on intermediation network models.

There is also a stream of papers that study intermediary problems in the context of supply chains (see, e.g., Belavina and Girotra 2012; Wu 2004; Nguyen et al. 2016), but these papers do not feature private information settings, unlike our models. Moreover, there are papers which consider the problem of successive monopolies in manufacturer-retailer settings with posted pricing schemes (see, e.g., Bresnahan and Reiss 1985; Lariviere and Porteus 2001; Perakis and Roels 2007). In the MSI model, we consider a more general class of mechanisms than posted pricing that allows intermediaries to elicit the private values of downstream agents and acquire the impression from upstream only if the downstream agents signal interest. Standard posted pricing mechanisms do not allow for this “contingent sale” feature, which is a prevalent in online advertising.

21.1.2.2 Online Advertising and Ad Exchanges

Mansour et al. (2012) give a brief synopsis of the auction employed by Google’s exchange, and discusses the role of intermediaries in this auction. Although this auction is not optimal or incentive compatible, the authors argue that the intermediaries’ incentives to misreport valuations is small in these auctions. The study by Ghosh et al. (2009) considers the advertiser’s perspective. In particular, Ghosh et al. (2009) study the design of a bidding agent who first explores the competing bids in the market, and then bids according to the observed empirical distribution. Balseiro

et al. (2015) and Gummadi et al. (2011) study the dynamic interactions among budget-constrained advertisers bidding in exchanges by proposing an approximate equilibrium concept. Iyer et al. (2014) characterize the bidding equilibrium among advertisers who learn about their private value over time in a repeated auction setting by using a mean-field approximation. Jiang et al. (2014) provide a simple bidding strategy, which does not require any statistical knowledge, for a single bidder with an average budget constraint. However these papers assume that advertisers bid directly in the exchange and do not consider the presence of intermediaries.

21.1.2.3 Mechanism Design with Budget Constraints

Finally, the OCI model considered in this chapter is related to a stream of papers that study mechanism design with financially constrained bidders (Laffont and Robert 1996; Che and Gale 1998, 2000; Maskin 2000). In these papers, efficient and optimal auctions are studied by modeling valuations as private and budgets as either private or public information. When the budgets are private, standard auction formats such as first-price and second-price auction are suboptimal and not revenue equivalent. Moreover, the problem in fact becomes a multi-dimensional mechanism design problem which is hard to solve in general. Pai and Vohra (2014) study the problem of selling one item in a setting with multiple budget-constrained buyers by using linear programming, and show that the optimal mechanism is implementable via an all-pay auction. Additionally, Chawla et al. (2011) show that the problem of selling one item to unit-demand buyers who are budget constrained can be reduced to an unconstrained problem with a small loss in performance. In multi-unit auctions with budget-constrained bidders, Borgs et al. (2005) and Bhattacharya et al. (2010) study approximation algorithms to design revenue-optimal incentive-compatible mechanisms. Brusco and Lopomo (2008, 2009) also study multi-item settings with budget constrained bidders by focusing on ascending auction formats. Differently from these papers, here we consider an intermediary who has no inherent value for the items sold and does not own the items at the moment of contracting, but instead needs to specify a mechanism to procure impressions as they arrive at the exchange. In addition, the main objective of this chapter is to understand the presence and profitability of such intermediaries.

21.2 Optimal Contracts for Intermediaries in Online Advertising

In this section, we provide the OCI model and the optimal mechanism design characterization. We consider a setting where an advertiser acquires impressions from an exchange over a fixed horizon (e.g., a week or a month) by either bidding directly at the exchange, or contracting with an *intermediary* to acquire impressions




Fig. 21.1 The advertiser can submit bids directly to the exchange, or through the intermediary. The exchange has other (exogenous) bidders as well

on her behalf. This contracting advertiser is referred to as “the advertiser” when clear from the context, to distinguish her from other exogenous advertisers. We assume that a fixed number, n , of impressions arrive at the exchange over the horizon.¹ In the following, we denote vectors using boldface as in \mathbf{x} , and the transpose of this vector by \mathbf{x}^T .

Exchange The exchange sells impressions to the contracting advertiser and other exogenous advertisers using a second-price auction with no reserve (see Fig. 21.1). Rather than explicitly modeling the preferences of the exogenous advertisers, we denote the exogenous advertisers’ maximum competing bid by d_i for impression i . We assume that $(d_i)_{i=1}^n$ are i.i.d. and drawn from the cumulative distribution function $F_d(\cdot)$, with the strictly positive probability density function $f_d(\cdot) > 0$ over the compact support $[0, \bar{D}]$.

For each arriving impression i , the exchange announces some user information in the form of an *attribute* vector, which may affect the value of the impression perceived by the advertiser. We denote by $\alpha_i \in \mathcal{A}$ an attribute vector which contains relevant information for the advertiser’s targeting criterion (such as geographical location, age group of the viewer, tastes and interests obtained from her browsing history). The space of attributes \mathcal{A} is a compact subset of the Euclidean space. We assume that the random variables $(\alpha_i)_{i=1}^n$ are i.i.d. with cumulative distribution function $F_\alpha(\cdot)$, and strictly positive density $f_\alpha(\cdot) > 0$. We next describe how the attribute vector α_i impacts the advertiser’s valuation of the impressions.

Advertiser The advertiser has a budget and a targeting criterion which we denote by $b \in \mathbb{R}_+$ and $\theta \in \Theta$, respectively. Therefore, the type of the advertiser is given by the pair $t = (b, \theta)$ that belongs to a set $\mathcal{T} \subseteq \mathbb{R}_+ \times \Theta$. The type t is the private information of the advertiser. The set \mathcal{T} is assumed to be finite, and we denote by $T \triangleq |\mathcal{T}|$ its cardinality. The advertiser’s type is $t \in \mathcal{T}$ with probability $p_t > 0$. With some abuse of notation we denote by b_t and θ_t the budget and targeting

¹Random number of impressions can be accommodated in our model by considering dummy arrivals that are valued at zero by the advertiser.

criterion of type t , respectively. The value of a type t advertiser for an impression with attributes α is given by $v_t(\alpha)$. Here the function $v_t : \mathcal{A} \rightarrow \mathbb{R}_+$ is bounded and continuous in the impression attributes $\alpha \in \mathcal{A}$ for every type $t \in \mathcal{T}$.

Outside Option The outside option for the advertiser corresponds to running a campaign on her own, and thus the value of the outside option is the maximum surplus the advertiser can get by participating in the exchange directly. If an advertiser with type $t = (b, \theta)$ pursues the outside option, she is constrained to the set of feasible non-anticipative policies (i.e., policies that map the history to bids) that satisfy the budget constraint for every sample path. This set of policies is denoted by \mathcal{Z}_t . Since the exchange runs a second-price auction, a feasible policy $\zeta \in \mathcal{Z}_t$ should satisfy the inequality:

$$\sum_{i=1}^n \mathbf{1}\{z_i^\zeta \geq d_i\} d_i \leq b_t \text{ (almost surely),}$$

where $\mathbf{1}\{\cdot\}$ is the indicator function, and $z_i^\zeta \in [0, \bar{D}]$ corresponds to the bid from policy ζ for the i th impression. The optimal expected surplus of an advertiser with type t is denoted by V_t , and obtained by solving the following optimal control problem:

$$V_t \triangleq \sup_{\zeta \in \mathcal{Z}_t} \mathbb{E}_{\alpha, \mathbf{d}} \left[\sum_{i=1}^n \mathbf{1}\{z_i^\zeta \geq d_i\} (v_t(\alpha_i) - d_i) \right], \quad (21.1)$$

where the expectation is taken with respect to the vector of impression attributes $\alpha = (\alpha_i)_{i=1}^n$ and maximum competing bids $\mathbf{d} = (d_i)_{i=1}^n$.

21.2.1 Mechanism Design Problem

In this section we focus on the problem of the intermediary whose objective is to run a campaign on behalf of the advertiser that can alternatively pursue her own campaign by directly participating in the auctions of the exchange. Since the advertiser's type (her budget and targeting criterion) is her private information, the intermediary's problem can naturally be formulated as a mechanism design problem. By the Revelation Principle, without loss of optimality, we can focus on direct mechanisms where the advertiser reports her type t (possibly nontruthfully), and the intermediary responds to this report by choosing a required payment x_t , and a dynamic bidding policy $\zeta_t \in \mathcal{Z}$ he commits to running at the exchange on behalf of the advertiser. Here, we denote by \mathcal{Z} the set of all non-anticipative (and

potentially randomized) dynamic bidding policies that have no budget restrictions.² Specifically, a mechanism (\mathbf{x}, ζ) for the intermediary consists of a vector of payments $\mathbf{x} = (x_t)_{t \in \mathcal{T}}$, and a vector of non-anticipative dynamic bidding policies associated with different types $t \in \mathcal{T}$, $\zeta = (\zeta_t)_{t \in \mathcal{T}}$.

Note that the set \mathcal{Z} of all non-anticipative bidding policies is a high-dimensional set which includes all functions mapping every possible history to a bid. Therefore, instead of searching for the optimal mechanism by optimizing directly over this set (which may be computationally intractable), we provide an alternative technique which relies on first characterizing the performance that can be achieved by policies in \mathcal{Z} . The performance of a policy $\zeta \in \mathcal{Z}$ can be measured by two metrics: (i) the intermediary’s total expected cost for running this policy

$$\mathcal{C}(\zeta) \triangleq \mathbb{E}_{\alpha, \mathbf{d}} \left[\sum_{i=1}^n \mathbf{1}\{z_i^\zeta \geq d_i\} d_i \right],$$

and (ii) the total expected value

$$\mathcal{W}_t(\zeta) = \mathbb{E}_{\alpha, \mathbf{d}} \left[\sum_{i=1}^n \mathbf{1}\{z_i^\zeta \geq d_i\} v_t(\alpha_i) \right]$$

an advertiser of type $t \in \mathcal{T}$ derives from the impressions acquired by this policy. We consider an optimal mechanism design problem where any performance level in this achievable performance space can be chosen by the intermediary.

Definition 1 The *achievable performance space* \mathcal{P} is given by the set of points $(c, \mathbf{w}) \subseteq \mathbb{R} \times \mathbb{R}^T$ such that there exists a policy $\zeta \in \mathcal{Z}$ satisfying $\mathcal{W}_t(\zeta) = w_t$ for all $t \in \mathcal{T}$ and $\mathcal{C}(\zeta) \leq c$.

Following the performance space idea, the optimal mechanism can be derived by optimizing over the costs and total expected values associated with feasible policies (as opposed to optimizing over the policies themselves). In comparison to the set \mathcal{Z} , the performance space has a much smaller dimension because it does not scale with the time horizon, thereby yielding a more tractable formulation. In addition, after the optimal performance levels associated with the optimal mechanism are determined, an optimal policy that achieves the chosen performance level can be retrieved via a *synthesis* procedure (described in Sect. 21.2.2.3).

We next introduce some notation. We use $c_{t'}$ to represent the total expected cost incurred by the intermediary when the advertiser reports her type as t' , and the intermediary uses the policy $\zeta_{t'}$ to bid on behalf of her at the exchange. If the true type of this advertiser is t , we denote her total expected value for the

²Unlike the advertiser, the intermediary does not have stringent financial constraints, thus we do not restrict the intermediary’s policies $\zeta \in \mathcal{Z}$ to satisfy any budget constraints (unlike the set of policies \mathcal{Z}_t that can be employed by the advertiser of type t).

impressions acquired by this policy by $w_{t',t}$, and the matrix of total expected values by $\mathbf{W} = (w_{t',t})_{t',t \in \mathcal{T}} \in \mathbb{R}^{T \times T}$. Note that the policy run for type t' may yield different expected values for advertisers of different types, i.e., for types t_1, t_2 advertisers, we may have $w_{t',t_1} \neq w_{t',t_2}$ because their targeting criteria may differ. Using this notation, the mechanism design problem of the intermediary can be stated as follows:

$$\max_{\mathbf{x}, \mathbf{c} \in \mathbb{R}^T, \mathbf{W} \in \mathbb{R}^{T \times T}} \sum_{t \in \mathcal{T}} p_t (x_t - c_t) \quad (21.2a)$$

$$\text{(OPT)} \quad \text{s. t.} \quad w_{t,t} - x_t \geq w_{t',t} - x_{t'}, \quad \forall t, t' : b_{t'} \leq b_t, \quad (21.2b)$$

$$w_{t,t} - x_t \geq V_t, \quad \forall t, \quad (21.2c)$$

$$x_t \leq b_t, \quad \forall t, \quad (21.2d)$$

$$(c_{t'}, (w_{t',t})_{t \in \mathcal{T}}) \in \mathcal{P}, \quad \forall t'. \quad (21.2e)$$

In (OPT), we maximize the expected profit of the intermediary, which is given by the quantity $\sum_{t \in \mathcal{T}} p_t (x_t - c_t)$. Here, the variable x_t represents the payment of the advertiser whose report is t and c_t is the cost of running the campaign. The *incentive compatibility (IC)* constraint Eq. (21.2b) ensures that the advertiser maximizes her payoff by reporting her type truthfully. Note that the surplus of a type t advertiser whose report is t' is expressed by the quantity $w_{t',t} - x_{t'}$. For the IC constraint, we restrict attention to the cases where the advertiser does not report a budget larger than her true budget without loss of generality.³ Recall that the advertiser could alternatively run her own campaign. Therefore, we guarantee that the mechanism delivers a utility at least equal to the outside option to the advertiser with the *individual rationality (IR)* constraint Eq. (21.2c). Moreover, the payment collected by the mechanism does not exceed the advertiser's budget due to the budget constraint Eq. (21.2d). The constraint Eq. (21.2e) guarantees that given an optimal solution of (OPT), the intermediary can find a bidding policy ζ_t that delivers the performance (to all types) as required by the optimal solution. Specifically, this constraint ensures that the performance of the mechanism of the intermediary lies in the achievable performance space \mathcal{P} , thereby implying that the structure of \mathcal{P} plays a key role in the solution of (OPT). Hence, we conclude this section by emphasizing that the performance space \mathcal{P} is convex and closed, which implies that (OPT) is a convex optimization problem.

Lemma 1 *The performance space \mathcal{P} is convex and closed.*

³Specifically, the intermediary can prevent the advertiser from overstating her budget by requiring her to make an upfront payment equal to the reported budget, and returning the amount $b_t - x_t$ at the end of the advertising campaign (see, e.g., Che and Gale 2000).

21.2.2 Optimal Mechanism Characterization

Because the performance space \mathcal{P} does not have closed-form description, solving (OPT) directly is algorithmically challenging. To overcome this difficulty, we introduce a duality-based approach for the optimal mechanism design problem which consists of dualizing some constraints of (OPT). This dual problem is a tractable convex minimization problem that can be solved efficiently.

21.2.2.1 Dual Problem

In the dual approach we dualize the IC, IR and budget constraints of (OPT), while optimizing over the performance space. To provide a characterization of the dual problem, we employ the concept of support function of the performance space. The support function $\phi : \mathbb{R} \times \mathbb{R}^T \rightarrow \mathbb{R}$ associated with the performance space \mathcal{P} for a point $(\mu, \boldsymbol{\lambda}) \in \mathbb{R} \times \mathbb{R}^T$ is given by

$$\phi(\mu, \boldsymbol{\lambda}) \triangleq \sup_{(c, \mathbf{w}) \in \mathcal{P}} \boldsymbol{\lambda}^\top \mathbf{w} - \mu c. \quad (21.3)$$

The support function ϕ gives the intercept of the supporting hyperplane of the performance space \mathcal{P} with normal $(-\mu, \boldsymbol{\lambda})$. The support function ϕ is convex because it is obtained as the point-wise supremum of linear functions over the performance space.

We simplify the derivation of the dual problem by momentarily writing the IC, IR and budget constraints in matrix form. In particular, we define $\mathbf{w}_t \triangleq (w_{t,t'})_{t' \in \mathcal{T}}$ as the t th row vector of the matrix $\mathbf{W} \in \mathbb{R}^{T \times T}$, and rewrite (OPT) as follows:

$$\max_{\mathbf{x}, \mathbf{c}, \mathbf{W}} \sum_{t \in \mathcal{T}} p_t(x_t - c_t) \quad (21.4a)$$

$$\text{s. t.} \quad \sum_{t \in \mathcal{T}} \mathbf{d}_t x_t - \mathbf{A}_t \mathbf{w}_t \leq \mathbf{e} \quad (21.4b)$$

$$(c_t, \mathbf{w}_t) \in \mathcal{P}, \quad \forall t. \quad (21.4c)$$

where $\mathbf{d}_t \in \mathbb{R}^M$, $\mathbf{A}_t \in \mathbb{R}^{M \times T}$ and $\mathbf{e} \in \mathbb{R}^M$ capture the coefficients of the $M \leq T^2 + 2T$ linear inequalities corresponding to the IC, IR and budget constraints.

Let $\boldsymbol{\lambda} \geq 0$ in \mathbb{R}^M be the Lagrange multiplier of constraints Eq. (21.4b). Using these multipliers, we obtain the convex dual problem (D):

$$\min_{\boldsymbol{\lambda} \in \mathbb{R}^M} \boldsymbol{\lambda}^\top \mathbf{e} + \sum_{t \in \mathcal{T}} \phi(p_t, \boldsymbol{\lambda}^\top \mathbf{A}_t) \quad (21.5a)$$

$$\text{(D) s. t.} \quad \boldsymbol{\lambda}^\top \mathbf{d}_t = p_t, \quad \forall t \quad (21.5b)$$

$$\boldsymbol{\lambda} \geq 0. \quad (21.5c)$$

In the following theorem, we show that strong duality holds for (OPT) and (D). This result is established in Balseiro and Candogan (2017) by first constructing a feasible solution for which the performance level associated with each type t belongs to the relative interior of \mathcal{P} . Using this constraint qualification with the known results from duality theory, the result follows.

Theorem 1 (OPT) admits an optimal solution. Additionally, strong duality holds, that is, the optimal objective value of (OPT) and (D) coincide.

We next show that the support function ϕ can be efficiently evaluated, and then discuss how to construct the optimal mechanism based on an optimal dual solution.

21.2.2.2 Support Function Characterization

Note that it is possible to characterize the support function more explicitly, by restating the expression in Eq. (21.3) using Definition 1:

$$\phi(\mu, \lambda) = \sup_{\zeta \in \mathcal{Z}, c \in \mathbb{R}, \mathbf{w} \in \mathbb{R}^T} \{ \lambda^\top \mathbf{w} - \mu c \text{ s.t. } \mathcal{C}(\zeta) \leq c, w_t = \mathcal{W}_t(\zeta) \}. \quad (21.6)$$

We next provide a closed-form expression for the support function as well as the optimal solution of Eq. (21.6). Let $\zeta^\phi(\mu, \lambda)$ be a policy that bids $z_i^{\zeta^\phi(\mu, \lambda)} = \sum_t \gamma_t v_t(\alpha_i) / \mu$ for impression i with attributes α_i , and define $c^\phi(\mu, \lambda) \triangleq \mathcal{C}(\zeta^\phi(\mu, \lambda))$ as the total expected cost of the policy and $w_t^\phi(\mu, \lambda) \triangleq \mathcal{W}_t(\zeta^\phi(\mu, \lambda))$ as the total expected value type $t \in \mathcal{T}$ has for the impressions acquired by this policy. In the following proposition we denote by $x^+ = \max(x, 0)$ the positive part of $x \in \mathbb{R}$.

Proposition 1 Suppose that $\mu > 0$. In Eq. (21.6), an optimal ζ is $\zeta^\phi(\mu, \lambda)$, and the unique optimal c and \mathbf{w} are given by $c^\phi(\mu, \lambda)$ and $\mathbf{w}^\phi(\mu, \lambda)$, respectively. More explicitly, the support function is given by

$$\phi(\mu, \lambda) = n \mathbb{E}_{\alpha, d} \left[\left(\sum_{t \in \mathcal{T}} \gamma_t v_t(\alpha) - \mu d \right)^+ \right],$$

the optimal total expected cost is $c^\phi(\mu, \lambda) = n \mathbb{E}_{\alpha, d} [d \mathbf{1}\{\sum_{t \in \mathcal{T}} \gamma_t v_t(\alpha) \geq \mu d\}]$, and the optimal total expected value for type $t \in \mathcal{T}$ is given by

$$w_t^\phi(\mu, \lambda) = n \mathbb{E}_{\alpha, d} \left[v_t(\alpha) \mathbf{1} \left\{ \sum_{t \in \mathcal{T}} \gamma_t v_t(\alpha) \geq \mu d \right\} \right].$$

Note that the support function can be expressed as a simple expectation (of a piecewise-linear function) over the impression attributes and the competing bid (Proposition 1). Moreover, the convex dual problem has a compact representation

(recall that the dual problem has polynomial size). Hence, it follows that the dual problem is tractable and its optimal solution can be obtained by using standard convex optimization algorithms.

21.2.2.3 Synthesis

In this section, we provide a procedure in which the optimal solution of the dual is used to “synthesize” the optimal contract of the intermediary by relying on strong duality. In particular, given an optimal solution λ^* of the dual problem (D), we first construct the optimal policy of the intermediary ζ_t^* and then obtain the corresponding performance (c_t^*, \mathbf{w}_t^*) for all $t \in \mathcal{T}$. We then solve a linear feasibility problem that involves the constructed performance levels in order to determine the upfront fees $\{x_t^*\}_{t \in \mathcal{T}}$ associated with the optimal contract. More formally, the steps of our synthesis procedure can be given as follows:

- Step 1. Determine an optimal solution λ^* of the dual problem (D), and set $\lambda_t^* = (\gamma_{t,t'}^*)_{t' \in \mathcal{T}} = (\lambda^*)^\top \mathbf{A}_t$ and $\mu_t^* = p_t$ for $t \in \mathcal{T}$.
- Step 2. Set the policy of type $t \in \mathcal{T}$ to $\zeta_t^* \triangleq \zeta^\phi(\mu_t^*, \lambda_t^*)$. The total expected cost and value for this policy are respectively given as $c_t^* \triangleq c^\phi(\mu_t^*, \lambda_t^*)$, and $\mathbf{w}_t^* \triangleq \mathbf{w}^\phi(\mu_t^*, \lambda_t^*)$ (see Proposition 1).
- Step 3. Determine upfront fees $\mathbf{x}^* = (x_t^*)_{t \in \mathcal{T}}$ by solving

$$\sum_{t \in \mathcal{T}} \mathbf{d}_t x_t^* - \mathbf{A}_t \mathbf{w}_t^* \leq \mathbf{e} \quad \perp \quad \lambda^* \geq 0,$$

where \perp indicates that for each entry, at least one of these inequalities should hold with equality.

Theorem 2 *The synthesis procedure yields an optimal solution $(\mathbf{x}^*, \mathbf{c}^*, \mathbf{W}^*)$ for (OPT) and an optimal mechanism (\mathbf{x}^*, ζ^*) for the intermediary.*

We note that the mechanism (\mathbf{x}^*, ζ^*) provided in Theorem 2 is relatively easy to implement. In this mechanism, the intermediary posts a menu of contracts (associated with different payments and policies), and invites the advertiser to choose the one she prefers. We next turn to the optimal policies of the intermediary and obtain further insights on these bidding policies.

21.2.2.4 Optimal Bidding Policy

Theorem 2 implies that the optimal bidding strategy has a surprisingly simple structure. In particular, assume that an impression with an attribute vector α arrives at the exchange, and the advertiser is of type t . Proposition 1 suggests that under the optimal policy the intermediary bids:

$$z_t^*(\alpha) = \frac{1}{p_t} \sum_{t' \in \mathcal{T}} \gamma_{t,t'}^* v_{t'}(\alpha). \quad (21.7)$$

It is possible to obtain further intuition on the bidding strategy of the intermediary. In particular, let $\lambda^* = (\lambda^{\text{IC}}, \lambda^{\text{IR}}, \lambda^{\text{B}})$, where $\lambda^{\text{IC}} = (\lambda_{t',t}^{\text{IC}})_{t',t}$, $\lambda^{\text{IR}} = (\lambda_t^{\text{IR}})_t$, $\lambda^{\text{B}} = (\lambda_t^{\text{B}})_t$, and $\lambda_{t',t}^{\text{IC}} \geq 0$, $\lambda_t^{\text{IR}} \geq 0$, $\lambda_t^{\text{B}} \geq 0$ respectively denote the optimal Lagrange multipliers associated with constraints (21.2b), (21.2c), and (21.2d) of (OPT). The next result characterizes the optimal bidding policy of the intermediary in terms of these multipliers.

Corollary 1 *Let $\lambda^* = (\lambda^{\text{IC}}, \lambda^{\text{IR}}, \lambda^{\text{B}})$ be a dual optimal solution of (D). The optimal bidding policy of the intermediary ζ_t^* for type $t \in \mathcal{T}$ is such that for an impression with attribute vector α it bids:*

$$z_t^*(\alpha) = \left(1 - \frac{\lambda_t^{\text{B}}}{p_t}\right) v_t(\alpha) + \frac{1}{p_t} \sum_{t' \in \mathcal{T}} \lambda_{t,t'}^{\text{IC}} (v_t(\alpha) - v_{t'}(\alpha)). \quad (21.8)$$

This corollary reveals that if the IC and budget constraints are not active (and hence $\lambda^{\text{IC}} = \lambda^{\text{B}} = 0$), the intermediary simply reports the true value $v_t(\alpha)$. For example, this case occurs when the advertiser has a large budget, and the impression distribution is uniform (since in this case the IC constraints are not binding, see Balseiro and Candogan (2017) for more details). Similarly, if the budget constraint is binding, but the IC constraints are not, then the intermediary simply shades the true value $v_t(\alpha)$ by $1 - \lambda_t^{\text{B}}/p_t$ to account for the advertiser's budget constraint.

When the IC constraints are active, the bidding strategy takes into account that type t may have incentive to misreport her type as t' (the dual multiplier of the corresponding incentive compatibility constraint is $\lambda_{t,t'}^{\text{IC}} \geq 0$). For instance, suppose that types t and t' have the same budget but t' has lower average valuations for the impressions than t . Type t' typically pays a lower amount for the impressions she acquires, and hence effectively has a less stringent budget constraint. Thus, the intermediary needs to bid more aggressively to match the outside option of type t' , which may give the type with higher valuations an incentive to impersonate the lower type. The term $(1/p_t) \sum_{t' \in \mathcal{T}} \lambda_{t,t'}^{\text{IC}} (v_t(\alpha) - v_{t'}(\alpha))$ in Eq. (21.8) ensures that the intermediary bids more aggressively for impressions that type t highly values when compared to other types (i.e., $v_t(\alpha) > v_{t'}(\alpha)$), thereby eliminating the incentive of type t to misreport her type.

21.2.3 Economic Insights

The presence of intermediation affects the online advertising market in several ways. As an immediate example, the optimal bidding policy of an advertiser participating in the exchange on her own has a complex dynamic shading structure, which is

obtained by solving a dynamic program. On the other hand, the policy associated with the optimal contract (see Corollary 1) of the intermediary has a stationary structure. Thus, from an operational perspective, the presence of the intermediary results in simpler bidding policies in the exchange. In addition to this impact on the bidding policies, we shed light on the other economic insights derived from the optimal mechanism of the intermediary by considering the following aspects.

21.2.3.1 Intermediation Profit and the Advertiser Surplus

Recall that the advertiser has the option of bidding directly in the exchange. Therefore, the intermediary should guarantee that the advertiser has incentive to accept the contract. In other words, the contract of the intermediary should deliver surplus to the advertiser at least as high as the surplus of the outside option. Despite providing this surplus to the advertiser, interestingly the intermediary still manages to profit. The reason behind this observation can be explained by the capability of the intermediary to deliver the surplus of the advertiser's outside option at a lower expected cost than the advertiser could achieve on her own. Specifically, the advertiser's bidding policies, when bidding directly in the exchange, have to satisfy her budget at every realization (of impression attributes and competing bids) whereas the intermediary has no budget constraints. Therefore, the advertiser's bidding policies might need to significantly shade her bid in each auction while the intermediary can implement bidding policies that exceed the upfront fee in some realizations. In other words, the absence of these financial constraints equips the intermediary with a richer set of policies, and allows him to deliver the same surplus to the advertiser at a lower cost; therefore making it profitable to intermediate.

21.2.3.2 Market Efficiency

Note that the optimal mechanism of the intermediary and the optimal bidding policy are provided for general valuation structures. Balseiro and Candogan (2017) provide further insights by focusing on a special case where advertisers' targeting criteria exhibit symmetry, e.g., uniformly distributed impression attributes. Under this assumption, the optimal mechanism can be characterized in closed form. Balseiro and Candogan (2017) also establish that the presence of the intermediary in the market not only allows him to profit, but also increases the surplus of the advertiser and market efficiency as given by the sum of the revenue of the exchange, the intermediary's profit, the contracting advertiser's surplus, and the exogenous bidders' surpluses.

21.3 Multi-stage Intermediation in Display Advertising

In this section, we consider a setting with multiple intermediaries positioned between an advertiser and an exchange (see Fig. 21.2). The exchange sells a unique indivisible impression via a second-price auction with no reserve in which exogenous advertisers might participate. Since we focus on the trade that occurs through the intermediation chain, as before, we model the highest bid of those exogenous advertisers by a random variable d with support $[0, \bar{D}]$. The advertiser seeks to purchase the impression from the exchange, and her value for the impression is captured by a random variable denoted by v .⁴ The realizations of v are drawn from the cumulative distribution function $G_v(\cdot)$, with the strictly positive probability density function $g_v(\cdot)$ over the support $\mathcal{V} \subseteq [0, \infty)$. We assume that the distribution of v is common knowledge, and its realizations are the private information of the advertiser. We also assume that the expected value of v is finite, i.e., $\mathbb{E}[v] < \infty$.

As opposed to the OCI model, the advertiser has no budget constraint, however, she is constrained to purchase impression from the exchange through a chain of $m \geq 1$ intermediaries. In Fig. 21.2, the intermediation chain is illustrated. We refer to the intermediaries closer to the advertiser as *downstream* intermediaries, and the intermediaries closer to the exchange as *upstream* intermediaries. The intermediaries have no value for the impression, thus they only profit in case of purchasing the impression from upstream and reselling to downstream. Therefore, a mechanism for an intermediary should (i) map reports from the downstream intermediary to an upstream bid to purchase the impression from upstream, and (ii) decide on an allocation and a payment to resell the impression to downstream.

For the set of available mechanisms, we focus on the set of second-price mechanisms \mathcal{M} where a mechanism $(r, Y) \in \mathcal{M}$ consists of a reserve price $r \in \mathbb{R}_+$ and a nondecreasing reporting function $Y : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. After receiving a downstream report x , intermediary j with mechanism (r_j, Y_j) reports $Y_j(x)$ if $x \geq r_j$, otherwise reports zero. The intermediary allocates the impression only if she wins it from upstream. The payment is determined as the minimum amount which guarantees winning of the downstream agent. Formally, given the exogenous bid $d \in [0, \bar{D}]$ at the exchange, we denote by \mathcal{W}_j the set of bids of intermediary I_j that guarantees winning, i.e.,




Fig. 21.2 A chain of intermediaries

⁴Note that in the OCI model we denote by $v_t(\alpha)$ the value of a type t advertiser for an impression with attributes α , thus the notation v_t represents a function. However, in the MSI model, the notation v represents a random variable which captures the value of the advertiser.

$$\mathscr{W}_j \triangleq \{x \geq 0 \mid x \geq r_{j+1}, Y_{j+1}(x) \geq r_{j+1}, \dots, Y_m \circ \dots \circ Y_{j+1}(x) \geq d\},$$

where Y_i and r_i respectively denote the reporting function and reserve of intermediary $I_{(i)}$ for $i = j + 1, \dots, m$, and \circ denotes the composition operator. The intermediary $I_{(j)}$ makes a payment to $I_{(j+1)}$ only in case of winning, and the payment amount is given by the smallest element of the winning reports \mathscr{W}_j , i.e., $\inf(\mathscr{W}_j)$. This payment rule is a natural extension of the second-price auction, which is commonly used in the context of display advertising. Therefore, the set of second-price mechanisms \mathscr{M} extends second-price auctions to a setting with intermediaries. We next provide some important properties of the set \mathscr{M} by the following lemma.

Lemma 2 *Suppose that intermediary $I_{(j)}$ selects her mechanism $(r_j, Y_j) \in \mathscr{M}$ for $j = 1, \dots, m$, and the exchange runs a second-price auction. Then,*

(a) *The composition of upstream mechanisms faced by the agent $I_{(j)}$ for $j = 0, \dots, m - 1$ (where $I_{(0)}$ represents the advertiser) is equivalent to a second-price auction with an exogenous random bid d^* and reserve price r^* given by*

$$d^* = Y_{j+1}^{-1} \circ \dots \circ Y_m^{-1}(d)$$

$$r^* = \max(r_{j+1}, Y_{j+1}^{-1}(r_{j+2}), Y_{j+1}^{-1} \circ Y_{j+2}^{-1}(r_{j+3}), \dots, Y_{j+1}^{-1} \circ \dots \circ Y_m^{-1}(0))$$

where we define the inverse of a reporting function Y as $Y^{-1}(x) = \inf\{\tilde{x} \geq 0 : Y(\tilde{x}) \geq x\}$.

(b) *Truthful bidding is an optimal strategy for the advertiser.*

The second item in this lemma implies that the truthful bidding is a best response for the advertiser regardless the mechanisms of the intermediaries. Therefore, we exclude the strategic behavior of the advertiser and focus on the game among intermediaries. We model the corresponding mechanism design game among intermediaries as a Stackelberg game where intermediaries move sequentially from upstream to downstream. In particular, the timing of the events is as follows:

1. The advertiser privately draws her value.
2. Intermediary $I_{(j)}$ determines a mechanism $(r_j, Y_j) \in \mathscr{M}$ after observing the mechanism of $I_{(j+1)}$ for $j = m, \dots, 1$.
3. The advertiser bids truthfully.
4. Intermediary $I_{(j)}$ learns the bid of $I_{(j-1)}$ and submits her own bid to $I_{(j+1)}$ according to her own mechanism, for $j = m, \dots, 1$.
5. The exchange $I_{(m+1)}$ runs a second-price auction and the impression is allocated either to exogenous bidders or to the advertiser through the intermediation chain.
6. Intermediary $I_{(j)}$ learns her payment to $I_{(j+1)}$ and charges a payment to $I_{(j-1)}$ for $j = m, \dots, 1$.

21.3.1 Equilibrium Characterization

We study the outcome of the strategic interaction among intermediaries by focusing on the subgame perfect equilibria (SPE) of the induced game as a solution concept and provide an SPE. Before stating our results, we provide some definitions that would be useful to characterize an SPE of this game. We first introduce the virtual value function and its inverse for a generic random variable X with a finite mean $\mathbb{E}[X] < \infty$, and c.d.f. $G_X(\cdot)$ and strictly positive p.d.f. $g_X(\cdot)$ over a nonnegative support $\mathcal{X} \subseteq [0, \infty)$.⁵ The virtual value function of X is given by

$$\phi_X(x) = x - \frac{1 - G_X(x)}{g_X(x)}$$

for $x \in \mathcal{X}$. Moreover, the inverse of the virtual value function is defined as $\phi_X^{-1}(x) \triangleq \inf\{\tilde{x} \in \mathcal{X} \mid \phi_X(\tilde{x}) \geq x\}$. We next define the *projected virtual value* function for random variables with strictly increasing virtual functions. This function projects the virtual value function to nonnegative reals, while extending its domain to \mathbb{R} .

Definition 2 Suppose X is an absolutely continuous random variable with a strictly increasing virtual value function $\phi_X(\cdot)$ and support \mathcal{X} . The projected virtual value function of X is given by

$$\psi_X(x) \begin{cases} \sup \mathcal{X} & x \geq \sup \mathcal{X} , \\ \phi_X(x) & z_X \leq x < \sup \mathcal{X} , \\ 0 & \text{otherwise,} \end{cases} \quad (21.9)$$

where the projection point is given by $z_X = \phi_X^{-1}(0)$. If the random variable X has an atom at zero and is absolutely continuous elsewhere in its support \mathcal{X} , then we define $\psi_X(\cdot)$ and z_X similarly by replacing $\phi_X(\cdot)$ in Eq. (21.9) with the virtual value $\phi_{X|X>0}(\cdot)$ of the strictly positive part of X , denoted by $X|X > 0$.⁶

In order to characterize an SPE of the game in the chain of intermediaries, we first focus on the mechanism design problem of a single intermediary positioned between the advertiser and the exchange, i.e., $m = 1$. The following lemma provides an optimal mechanism for this case.

⁵The advertiser's value v satisfies these requirements.

⁶Note that $\phi_X(x) = \phi_{X|X>0}(x)$ for $x \in \mathcal{X} \setminus \{0\}$. This can be seen by using the definition of the virtual value function and noting that the conditional random variable $X|X > 0$ has c.d.f. $G_{X|X>0}(x) = (G_X(x) - G_X(0))/(1 - G_X(0))$, and p.d.f. $g_{X|X>0}(x) = g_X(x)/(1 - G_X(0))$. Thus, focusing on the virtual value of $X|X > 0$ as opposed to X , excludes the atom at zero, without impacting the (projected) virtual values elsewhere.

Lemma 3 *Suppose that there exists a single intermediary between the advertiser and the exchange, i.e., $m = 1$. Then, an optimal mechanism $(r, Y) \in \mathcal{M}$ for the intermediary is given by*

$$Y(x) = \psi_v(x),$$

$$r = z_v.$$

When choosing her bidding strategy for the upstream auction, an intermediary needs to optimally tradeoff the probability of winning the impression in the auction of the exchange with the cost incurred when winning the impression (both of which increase with the bid of the intermediary). This lemma reveals that the intermediary's optimal bidding strategy takes a simple structure: the intermediary first determines the advertiser's virtual value function, and then bids at the exchange the projected virtual value of the report from downstream.

In the case of multiple intermediaries, there are two factors which can affect the strategy of an intermediary: (i) the mechanisms chosen by upstream intermediaries, and (ii) the reaction of downstream intermediaries. For the first factor, Lemma 2 suggests that in settings with multiple intermediaries, the upstream mechanism that an intermediary faces can equivalently be represented by a second-price auction. Therefore, in light of Lemma 3, an intermediary along the chain is not influenced by the upstream decisions and, in turn, her actions do not influence downstream mechanisms. For the second factor, note that the reports observed by an intermediary $I_{(j)}$ can be obtained by composing the reporting functions of intermediaries $I_{(j-1)}, \dots, I_{(1)}$ with the report of the advertiser. Specifically, each intermediary in the chain can be thought as a single intermediary which connects a "downstream agent" with "anticipated reports" induced by the optimal downstream mechanisms along the equilibrium path to an upstream second-price auction. This observation allows for characterizing the optimal mechanisms of intermediaries recursively via Lemma 3, starting with the downstream intermediaries whenever the downstream reports satisfy the regularity conditions as the advertiser's value does. Therefore, we first formally define the anticipated reports and provide the regularity assumption, and characterize an SPE of the game between multiple intermediaries.

Definition 3 The anticipated report of the advertiser is $W_0 = v$, and the anticipated report of intermediary $I_{(j)}$ is

$$W_j = \psi_{W_{j-1}}(W_{j-1}).$$

Note that the anticipated report W_j coincide with the report of intermediary $I_{(j)}$ to the upstream mechanism if all intermediaries use the projected virtual value functions of the downstream bids as reporting functions.

Assumption 3 *The anticipated report W_j of intermediary $I_{(j)}$ is well-defined and has finite expected values for $j = 1, \dots, m$. Moreover, the anticipated reports have strictly increasing virtual values, i.e., $\phi_{W_j}(\cdot)$ (or $\phi_{W_j|W_j>0}(\cdot)$ if W_j has an atom at zero) is strictly increasing for $j = 1, \dots, m$.*

The following theorem provides an SPE under Assumption 3.⁷

Theorem 4 *Suppose that Assumption 3 holds, and consider the mechanisms given by*

$$Y_j^*(x) = \psi_{W_{j-1}}(x),$$

$$r_j^* = z_{W_{j-1}},$$

for $j = 1, \dots, m$. Then, the strategy profile where the mechanism (r_j^*, Y_j^*) is chosen by intermediary $I_{(j)}$ for $j = 1, \dots, m$ constitutes an SPE.

21.3.2 Economic Insights

Using the SPE provided in Theorem 4, we numerically explore the impact of the advertiser's value distribution on the reports and profits of intermediaries in different positions. Finally, we compare our results with the double-marginalization literature.

Intermediaries' Bids Since the virtual value function lies below the 45 degree line, Theorem 4 shows that the intermediaries always shade their bids (i.e., the report functions satisfy $Y_j^*(x) \leq x$). Although bid shading is always present, our numerical studies illustrate that it takes a different structure depending on the advertiser's value distribution. Figure 21.3 indicates that as a result of bid shading in long intermediation chains, unless the value of the buyer is significantly large, the bid submitted by the intermediaries to the auction of the exchange can be equal to zero. In such cases, even when it is profitable for the buyer, the intermediation chain does not allocate the impression to the advertiser, thereby causing inefficiency.

Intermediaries' Profits We next study the impact of an intermediaries position in the chain on her profits. On one hand, downstream intermediaries closer to the advertiser receive higher bids. On the other hand, upstream intermediaries closer to the exchange incur lower payments to acquire the impression. The total contribution of these opposing effects is indeterminate, and the impact of the position on profits depends on the distribution of the advertiser's value.

⁷Balseiro et al. (2017) show that this assumption holds for Generalized Pareto Distributions, a large family of distributions including uniform, exponential and Pareto distributions.

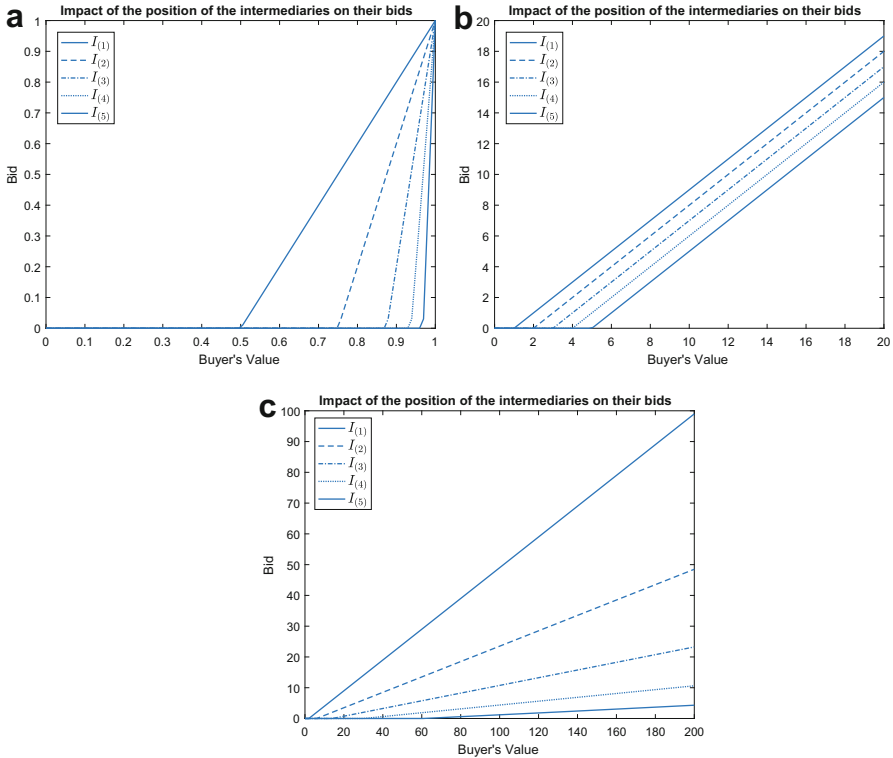


Fig. 21.3 This figure plots the bids of the intermediaries for different realization of the advertiser’s value v in a market with $m = 5$ intermediaries when the value distribution is uniform, exponential and shifted Pareto **(a)** Uniform distribution ($\mathcal{V} = [0, 1], \mathbb{E}[v] = 1/2$), **(b)** Exponential distribution ($\mathcal{V} = [0, \infty), \mathbb{E}[v] = 1$), **(c)** S. Pareto distribution ($\mathcal{V} = [0, \infty), \mathbb{E}[v] = 2/3$)

For example, when values are exponentially distributed the expected profit of each intermediary is the same, while for the uniform and shifted Pareto distributions the profits depend on the position of the intermediary in the chain. In particular, Fig. 21.4 plots the impact of the intermediary’s position on her profits conditional on winning the impression in a market with $m = 5$ intermediaries when the distribution of values is uniform, exponential and shifted Pareto. This figure shows that for uniform and Pareto distributions, profits as a function of an intermediary’s position in the chain, exhibit different trends. When the advertiser has a heavy-tailed value distribution, such as the shifted Pareto distribution, the downstream intermediaries who are closer to the advertiser have higher profits. Intuitively, this result stems from the fact that for such distributions, with significant probability the value of the advertiser for the impression is large, and hence the intermediaries that are closer to the advertiser can claim significant profits. Conversely, when the advertisers value distribution has a light-tail or a bounded support, the intermediaries find it more profitable to be closer to the exchange, due to lower costs of acquiring impressions.

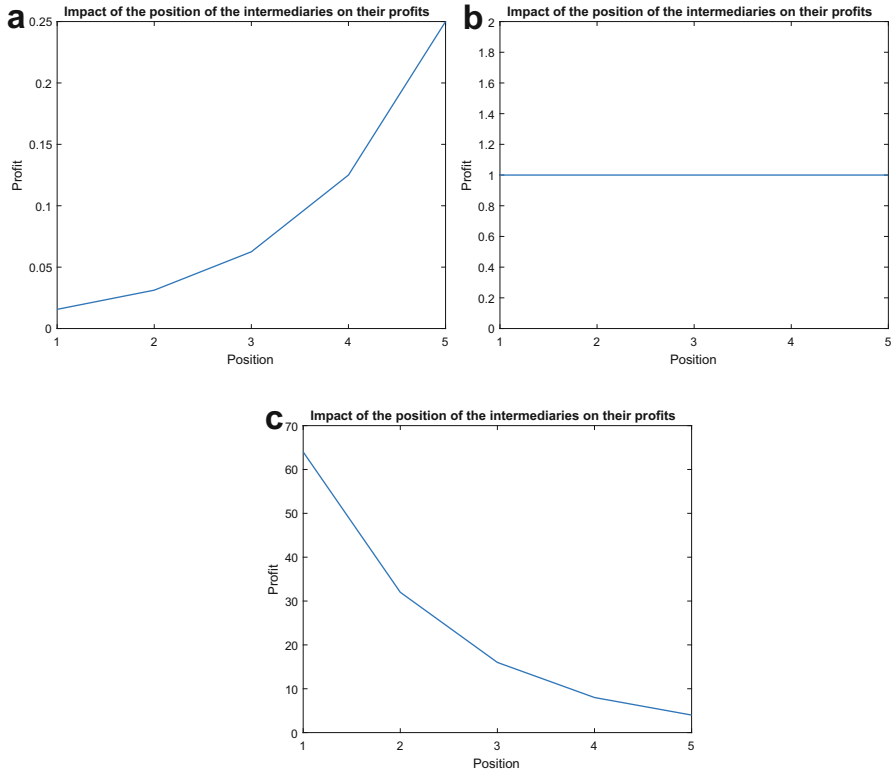


Fig. 21.4 This figure plots the profits of intermediaries for different positions in a market with $m = 5$ intermediaries when the value distribution is uniform, exponential and shifted Pareto ((**a**) Uniform distribution ($\mathcal{V} = [0, 1]$, $\mathbb{E}[v] = 1/2$), (**b**) Exponential distribution ($\mathcal{V} = [0, \infty)$, $\mathbb{E}[v] = 1$), (**c**) S. Pareto distribution ($\mathcal{V} = [0, \infty)$, $\mathbb{E}[v] = 2/3$))

This result suggests that depending on the advertiser’s value distribution for the impressions, intermediaries may prefer to participate in different stages of the intermediation process.

Comparison with Double-Marginalization Literature Our work is closely related to the double-marginalization literature (see, e.g., Tirole 1988), and some of the insights from this literature translate to our settings. In the basic double-marginalization framework, a manufacturer supplies a good to a single downstream retailer, who resells the good as a monopolist. Compared to the vertically integrated industry, the theory predicts that in the case of double-marginalization the price paid by the consumers is higher, industry profits, and the overall market efficiency are lower.

In our setting the first-best corresponds to the case when the advertiser bids directly in the exchange’s auction. Compared to the first-best, it is not hard to see that as the number of intermediaries increases, the expected surplus of the advertiser,

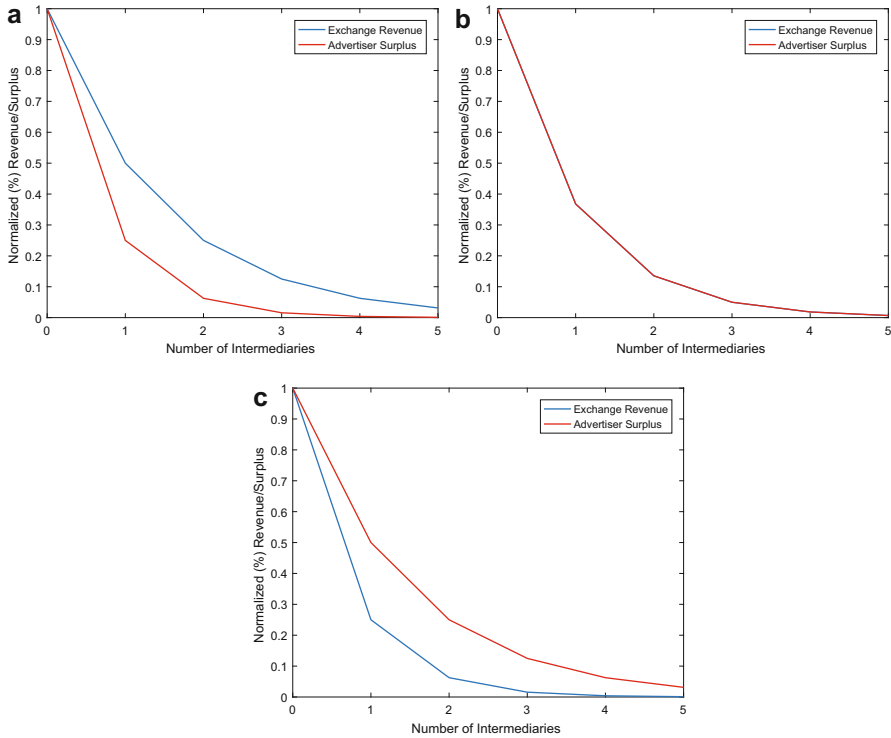


Fig. 21.5 Impact of the total number of intermediaries (m) on the expected surplus of the advertiser ($I_{(0)}$), and the expected revenue of the exchange ($I_{(m+1)}$). All results are relative to the first-best, which corresponds to the case when the advertiser bids directly in the exchange’s mechanism ($m = 0$). In these plots, the value of the exogenous bid d is constant and equal to $\mathbb{E}[v]$ ((a) Uniform distribution ($\mathcal{V} = [0, 1]$, $\mathbb{E}[v] = 1/2$), (b) Exponential distribution ($\mathcal{V} = [0, \infty)$, $\mathbb{E}[v] = 1$), (c) S. Pareto distribution ($\mathcal{V} = [0, \infty)$, $\mathbb{E}[v] = 2/3$))

revenue of the exchange, are affected negatively. This is a direct consequence of the bid shading behavior exhibited by the intermediaries. Figure 21.5 plots the expected surplus of the advertisers and the expected revenue of the exchange relative to that of first-best. This figure suggests that even one intermediary can significantly decrease the exchange’s revenue and the surplus of the advertiser, while multiple intermediaries can quickly decrease these quantities to almost zero. This result is aligned with the double marginalization literature. Our analysis extends these double-marginalization insights by quantifying the impact of the distribution of values and the position in the intermediation chain on an agent’s profit. In particular, when the value distribution is uniform (shifted Pareto) the advertiser (the exchange) suffers more than the exchange (the advertiser) from intermediation. On the other hand, the exchange and the advertiser are equivalently affected when the value distribution is exponential. These observations are aligned with the optimal intermediation position in multi-stage intermediation settings identified in our analysis.

21.4 Concluding Remarks

We study two models which shed light on the problems related to intermediation in online advertising. In the first part, we characterize the optimal mechanism for an intermediary which offers a contract to an advertiser with a private budget and targeting criterion to acquire impression on her behalf. We show that the presence of the intermediary does not harm the advertiser surplus, even further leads to a simpler bidding policy. For this model, Balseiro and Candogan (2017) also show that the profits of the intermediary are maximized for markets where budgets of advertisers are neither exceedingly small nor exceedingly large by using a combination of theoretical results and numerical experiments. In the second part, we provide a game theoretic model to understand the strategic interaction between intermediaries organized in a chain network, and derive economic insights via numerical analyses. As an immediate extension of this model, more general network structures can be considered in addition to the chain network. To this end, Balseiro et al. (2017) study symmetric tree networks, and formally establish that the results shown in the second part of this chapter, such as the impact of the value distribution on the most profitable position in a network, are generalized to those networks. Moreover, they also analyze the incentives of intermediaries to merge horizontally (within the same tier) and vertically (across different tiers). Although the exchange considered herein is assumed to have no reserve price, Balseiro et al. (2017) also model the strategic behavior of the exchange, and show that the intermediation network structure plays a key role in the profit of the exchange.

References

- Balseiro SR, Candogan O (2017) Optimal contracts for intermediaries in online advertising. *Oper Res* 65(4):878–896
- Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: approximations and design. *Manag Sci* 61(4):864–884
- Balseiro SR, Candogan O, Gurkan H (2017) Multi-stage intermediation in display advertising. Working paper
- Belavina E, Girotra K (2012) The relational advantages of intermediation. *Manag Sci* 58(9):1614–1631
- Bhattacharya S, Goel G, Gollapudi S, Munagala K (2010) Budget constrained auctions with heterogeneous items. In: Proceedings of the 42nd ACM symposium on theory of computing, STOC'10. ACM, New York, pp 379–388
- Borgs C, Chayes J, Immorlica N, Mahdian M, Saberi A (2005) Multi-unit auctions with budget-constrained bidders. In: Proceedings of the 6th ACM conference on electronic commerce, EC'05. ACM, New York, pp 44–51
- Bresnahan TF, Reiss PC (1985) Dealer and manufacturer margins. *RAND J Econ* 16:253–268
- Brusco S, Lopomo G (2008) Budget constraints and demand reduction in simultaneous ascending-bid auctions. *J Ind Econ* 56(1):113–142
- Brusco S, Lopomo G (2009) Simultaneous ascending auctions with complementarities and known budget constraints. *Econ Theory* 38(1):105–124
- Chawla S, Malec DL, Malekian A (2011) Bayesian mechanism design for budget-constrained agents. In: Proceedings of the 12th ACM conference on electronic commerce, EC'11. ACM, New York, pp 253–262

- Che YK, Gale I (1998) Standard auctions with financially constrained bidders. *Rev Econ Stud* 65(1):1–21
- Che YK, Gale I (2000) The optimal mechanism for selling to a budget-constrained buyer. *J Econ Theory* 92(2):198–233
- Condorelli D, Galeotti A (2016) Strategic models of intermediation networks. In: Bramoulle Y, Galeotti A, Rogers BW (eds) *The Oxford handbook of the economics of networks*. Oxford University Press, New York
- Feldman J, Mirrokni V, Muthukrishnan S, Pai MM (2010) Auctions with intermediaries: extended abstract. In: *Proceedings of the 11th ACM conference on electronic commerce, EC'10*. ACM, New York, pp 23–32
- Ghosh A, Rubinfeld BI, Vassilvitskii S, Zinkevich M (2009) Adaptive bidding for display advertising. In: *Proceedings of the 18th international conference on World Wide Web, WWW'09*. ACM, New York, pp 251–260
- Gummadi R, Key PB, Proutiere A (2011) Optimal bidding strategies in dynamic auctions with budget constraints. In: *49th annual Allerton conference on communication, control, and computing (Allerton)*. IEEE, Piscataway, p 588
- Internet Advertising Bureau (2016) Internet advertising revenue report, 2016 full year results. PricewaterhouseCoopers. Technical report. Available at https://www.iab.com/wp-content/uploads/2016/04/IAB_Internet_Advertising_Revenue_Report_FY_2016.pdf. Accessed 6 Mar 2018
- Iyer K, Johari R, Sundararajan M (2014) Mean field equilibria of dynamic auctions with learning. *Manag Sci* 60(12):2949–2970
- Jiang C, Beck CL, Srikant R (2014) Bidding with limited statistical knowledge in online auctions. *SIGMETRICS Perform Eval Rev* 41(4):38–41
- Laffont JJ, Robert J (1996) Optimal auction with financially constrained buyers. *Econ Lett* 52(2):181–186
- Lariviere MA, Porteus EL (2001) Selling to the newsvendor: an analysis of price-only contracts. *Manuf Serv Oper Manag* 3(4):293–305
- Loertscher S, Niedermayer A (2007) When is seller price setting with linear fees optimal for intermediaries? Technical report, Department of Economics, Universität Bern, Discussion papers
- Loertscher S, Niedermayer AF (2012) Fee-setting mechanisms: on optimal pricing by intermediaries and indirect taxation. Available at SSRN 2172386
- Manea M (2018) Intermediation and resale in networks. *J Polit Econ* 126(3):1250–1301
- Mansour Y, Muthukrishnan S, Nisan N (2012) Doubleclick ad exchange auction. <https://arxiv.org/pdf/1204.0535v1.pdf>
- Maskin ES (2000) Auctions, development, and privatization: efficient auctions with liquidity-constrained buyers. *Eur Econ Rev* 44(4–6):667–681
- Myerson RB, Satterthwaite MA (1983) Efficient mechanisms for bilateral trading. *J Econ Theory* 29(2):265–281
- Nguyen T, Subramanian V, Berry R (2016) Delay in trade networks. *Oper Res* 64(3):646–661
- Niazadeh R, Yuan Y, Kleinberg R (2014) Simple and near-optimal mechanisms for market intermediation. In: *Web and Internet economics*. Springer, Cham, pp 386–399
- Pai MM, Vohra R (2014) Optimal auctions with financially constrained bidders. *J Econ Theory* 150(0):383–425
- Perakis G, Roels G (2007) The price of anarchy in supply chains: quantifying the efficiency of price-only contracts. *Manag Sci* 53(8):1249–1268
- Stavrogianis LC, Gerding EH, Polukarov M (2013) Competing intermediary auctions. In: *Proceedings of the 2013 international conference on autonomous agents and multi-agent systems*, St. Paul, pp 667–674
- Tirole J (1988) *The theory of industrial organization*, 1st edn. The MIT Press, Cambridge/London
- Wu SD (2004) Supply chain intermediation: a bargaining theoretic framework. In: Simchi-Levi D, Wu SD, Shen Z-J (eds) *Handbook of quantitative supply chain analysis*. Springer, New York, pp 67–115