Optimal Recommendation Mechanisms

Changhwa Lee*

March 16, 2026

Abstract

An intermediary monetizes data by selling recommendations to sellers that influence consumer choice. The intermediary cannot control the consumer's decision, so recommendations must remain sufficiently informative for the consumer to follow them, while sellers are partially informed about their willingness to pay. I characterize the optimal recommendation mechanism, which recommends the option with the highest adjusted virtual willingness to pay. It can be implemented in weakly dominant strategies as a weighted second-maximum willingness-to-pay auction. The auction generalizes the second-price auction to environments with persuasion and partial information while preserving tight bid bounds on equilibrium payments.

1 Introduction

Digital platforms collect vast amounts of data about consumer preferences and behavior. As intermediaries between consumers and sellers, they frequently monetize this data by selling recommendations that influence consumer choice while charging firms for favorable placement. Search engines rank sponsored links, e-commerce platforms highlight promoted products, and AI assistants increasingly feature sponsored recommendations. These recommendation markets have become a central channel through

*Department of Economics, University of Bristol. Email: changhwalee.cl@gmail.com. I am deeply indebted to Rakesh Vohra, George Mailath and Andrew Postlewaite for their guidance and support. I thank In-Uck Park, Kevin He, Erik Madsen, and seminar participants at various universities and conferences for helpful comments. An earlier version of this paper circulated as *Optimal Recommender System Design* (Lee, 2023). All errors are my own.

which platforms monetize data, rather than selling it directly, and have attracted increasing regulatory scrutiny.[1]

Selling recommendations creates a distinctive mechanism design problem. The intermediary observes match information between the consumer and products that neither the consumer nor sellers observe. This information allows the intermediary to influence consumer choice through recommendations but not to control it: after observing a recommendation, the consumer chooses the option with the highest expected value. Recommendations must therefore remain sufficiently informative for the consumer to follow them, imposing a persuasion constraint on revenue extraction. Meanwhile, sellers are willing to pay to be recommended, but their willingness to pay depends on match information observed only by the intermediary, leaving them only partially informed about it. The intermediary must therefore elicit sellers' private information and combine it with match data to infer sellers' willingness to pay while maintaining credible recommendations. The mechanism design problem therefore combines persuasion, screening, and information aggregation.

A natural approach is to auction recommendations, but standard auction formats fail. In standard auctions, recommendations depend solely on bids. Because sellers do not observe match information, bids cannot reflect it. Recommendations therefore fail to aggregate information about match quality and are ignored by the consumer. When persuasion fails, sellers have no incentive to bid. Bidding zero weakly dominates, yielding zero revenue.

This paper studies the optimal design of recommendation mechanisms. An intermediary privately observes match values between a consumer and several products. Sellers therefore have only partial information about their willingness to pay. The intermediary conditions recommendations on match values and charges sellers. After observing a recommendation, the consumer updates beliefs about match quality and chooses the option with the highest expected value.

The first main result characterizes the revenue-maximizing recommendation mechanism. Absent persuasion constraints, the intermediary would recommend the option with the highest virtual willingness to pay, reflecting the marginal revenue generated by each seller. However, recommendations must remain sufficiently informative for

[1]For example, advertising constitutes over 95% of Meta's revenues and over 75% of Alphabet's revenues. See Alphabet Inc., *Form 10-K Annual Report* (2024), at 30, https://s206.q4cdn.com/479360582/files/doc_financials/2024/q4/goog-10-k-2024.pdf.

the consumer to follow them. The optimal rule therefore recommends the option with the highest non-negative *adjusted virtual willingness to pay*, where the adjustment ensures that recommendations remain informative.

Characterizing this optimal rule involves reformulating the intermediary's problem as a Bayesian persuasion problem. The resulting persuasion problem features a continuous, high-dimensional state space, multiple actions, and state-dependent sender payoffs, where standard tools are less tractable. I introduce a tractable monotonicity condition, *swap-monotonicity*, that is sufficient for obedience across symmetric products and admits a partial converse within symmetric recommendation rules, and allows a direct characterization of optimal mechanisms in multi-option environments.

The second main result establishes implementation and robustness. The optimal mechanism can be implemented in weakly dominant strategies by a *weighted second maximum willingness to pay auction (wSMWA)*. Sellers submit bids representing their maximum willingness to pay, which the intermediary combines with its information to infer sellers' valuations and applies second-price logic. Among a broad class of auction rules implementing the optimal mechanism, the wSMWA uniquely satisfies a tight upper-bound payment property, equilibrium payments are tightly bounded above by bids, as in the standard second-price auction. In this sense, the wSMWA extends the second-price auction to environments with persuasion and partial information.

This implementation provides a mechanism design foundation for bidding formats widely used in digital advertising platforms. In practice, advertisers submit bids reflecting their maximum willingness to pay for an impression rather than their exact valuations. Platforms then combine these bids with proprietary match-quality data to determine rankings and payments. The analysis shows that such structures arise naturally when platforms possess information unavailable to sellers.

Furthermore, the wSMWA is robust across all ranges of informational and behavioral assumptions. I characterize the optimal auction across environments that vary along three dimensions: whether the consumer's choice is allocated or implemented through persuasion, whether sellers are fully or partially informed about their valuations, and whether the intermediary observes match quality. The wSMWA is optimal in all environments and reduces to the (weighted) second-price auction whenever that format is optimal, thereby robustly generalizing them regardless of whether persuasion or partial information is present. Table 1 summarizes the taxonomy.

Finally, I study how changes in information access affect the optimal mechanism.

Consumer	Seller	Intermediary	SPA	wSPA	wSMWA
Allocative	Fully	Uninformed	O	O	O
		Informed	X	O	O
	Partially	Informed	X	X	O
Persuasive	Fully	Uninformed	X	O	O
		Informed	X	O	O
	Partially	Informed	X	X	O

Table 1: Optimal auction formats across relevant environments. SPA, wSPA, and wSMWA denote the second-price auction, weighted second-price auction, and weighted second maximum willingness to pay auction, respectively. O and X denote, respectively, that the auction (weakly dominantly) implements an optimal mechanism or fails to do so in the given environment.

Allowing the intermediary to sell or disclose its match information to sellers, or allowing sellers to acquire such information themselves, does not change the optimal mechanism, because recommendations and payments already condition on this information. In contrast, restricting the intermediary's access to match data and sellers' information reduces the informativeness of recommendations and harms consumers with poor outside options. These results connect to broader literatures on the sale of information and information in auctions, and to policy debates on data access and information sharing in digital platforms.

The rest of the paper is organized as follows. In the remaining part of this section, I discuss related literature. Section 2 introduces the model. In Section 3 characterizes the optimal recommendation mechanism. Section 4 studies implementation and robustness. Section 5 analyzes information access and its policy implications. Section 6 provides examples. The appendix contains omitted proofs.

1.1 Related Literature

Sales of Information. This paper is most closely related to the literature on the sales of information. Much of this literature on information sales focuses on *direct sales of information*, in which information itself is the object of trade (Admati and Pfleiderer, 1990), including models of selling experiments (Bergemann, Bonatti, and

Smolin, 2018), statistics (Segura-Rodriguez, 2021), market segments (Yang, 2021) and buyer's valuations (Evans and Park (2026), Ichihashi and Smolin (2025b)).

By contrast, this paper studies *indirect sales of information*, where the seller monetizes information through actions or allocations informed by it rather than directly transferring it. Admati and Pfleiderer (1990) analyze indirect monetization through financial portfolios. Closely related are Yang (2024), who studies an intermediary designing recommendation, transfer, and pricing rules for a single product, Bergemann and Bonatti (2024), who analyze how digital platforms allocate recommendation among competing sellers using consumer data, and Ichihashi and Smolin (2025a), who study the design of buyer-optimal recommendation algorithms. These papers examine how recommendation or allocation rules shape market outcomes. I instead study a multi-product mechanism design framework with exogenous prices.

A broader literature studies platform steering (Inderst and Ottaviani (2012), Teh and Wright (2022), Nocke and Rey (2024)), where steering works through commissions, rankings, or recommendation decisions. Here, steering is purely informational.

Sponsored Auctions. The paper also contributes to the sponsored-auctions literature. Classic analyses assume allocative consumer choice and fully informed advertisers, yielding generalized second-price logic (Edelman, Ostrovsky, and Schwarz (2007), Varian (2007)). Subsequent work allows auction outcomes to convey information that affects consumer choice, including environments in which sellers are better informed than the intermediary about match quality (Athey and Ellison, 2011), symmetrically informed (Gomes, 2014), or less informed (Bergemann, Bonatti, and Wu, 2025). Closest to this paper is Gomes (2014) who shows that a weighted second-price auction is optimal when sellers are symmetrically informed. I contribute by studying a mechanism-design framework in which the intermediary is better informed than sellers about consumer demand and sellers are only partially informed, and by deriving an auction format that nests the standard and weighted second-price auctions while remaining optimal across these environments.

Mechanism Design. More broadly, this paper relates to mechanism design in environments where the designer cannot fully control outcomes. In standard mechanism design, allocations are enforceable and incentive constraints operate directly through assignment and transfers (Myerson, 1981). Here, the intermediary cannot enforce the consumer's choice and must instead rely on persuasion: a mechanism is feasible only if recommendations are sufficiently informative to induce obedience. As a result, seller discipline operates through induced choice probabilities rather than direct allocation. This connects the paper to work on mechanism design with endogenous actions following messages (Myerson (1982), Myerson (1983)) and to environments in which the designer interacts with a subsequent stage that it cannot control (Dworczak, 2020).

Bayesian Persuasion. Finally, the paper contributes to the Bayesian persuasion literature (Rayo and Segal (2010), Kamenica and Gentzkow (2011), Bergemann and Morris (2017)). A large body of work studies persuasion using concavification (Kamenica and Gentzkow, 2011), convex analysis (Gentzkow and Kamenica, 2016), duality (Kolotilin (2018), Galperti and Perego (2018), Dworczak and Kolotilin (2019), Dworczak and Martini (2019)), optimal transport (Kolotilin, Corrao, and Wolitzky, 2025) and recursive methods (Smolin and Yamashita, 2025). The present paper considers a continuous, high-dimensional state space with multiple actions and state-dependent sender payoffs, where these standard tools do not readily apply. I identify a tractable condition, swap-monotonicity, that is sufficient for obedience in symmetric multi-option environments, and use it to characterize the optimal solution.

2 Model

There is a representative consumer (he), an intermediary (it) and $N \geq 2$ sellers (she), each offering a product. The intermediary has access to the consumer and seller data, from which it derives *match values* $\boldsymbol{v} = (v_1, \cdots, v_N)$ between the consumer and products, indexed by $i \in \{1, \cdots, N\} = \mathcal{N}$. Each match value v_i is independently drawn from a common cumulative distribution function $F(.)$ that has a compact support $\mathcal{V} \subset \mathbb{R}^1$ with $\underline{v} = \min \mathcal{V} < \max \mathcal{V} = \overline{v}$. Each seller $i \in \mathcal{N}$ has a *private type* θ_i that is independently drawn from a common cumulative distribution function $G(.)$ with support $[\underline{\theta}, \overline{\theta}]$. The distribution G admits a positive, continuous density $g(.)$ on its support and satisfies the monotone hazard rate condition.

The consumer chooses one of the products or an outside option, but does not know either the match values or the sellers' private types. If the consumer selects product i, he receives utility v_i. If he chooses the outside option, denoted by $i = 0$, he receives a commonly known *outside option value* $v_0 \in (\underline{v}, \overline{v})$.

If her product is selected, seller i earns payoff $w(v_i, \theta_i)$; otherwise her payoff is zero. The function w is continuous in (v, θ), non-decreasing in v, and is continuously differentiable in θ with $w_\theta > 0$. Thus, $w(v_i, \theta_i)$ represents seller i's *willingness to pay* for inducing the consumer to choose her product. If seller i is chosen with probability r_i and pays transfer t_i, her expected payoff is

$$w(v_i, \theta_i)r_i - t_i.$$

Sellers observe their own types θ_i but not match values v_i.

One interpretation of $w(v_i, \theta_i)$ is sponsored search. A consumer enters a keyword that is publicly observed by sellers, while the platform privately observes how well each seller's link matches the consumer. This match quality may be inferred from user data such as search history or past behavior and affects downstream outcomes such as engagement or purchase probability. Sellers differ in their underlying profitability, summarized by a privately known type θ_i. Conditional on being chosen, higher match values increase expected profits, so $w(v_i, \theta_i)$ is increasing in v_i. Sellers observe their own profitability but not the realized match quality and are therefore only partially informed about their willingness to pay.

Define the *virtual willingness to pay* φ by

$$\varphi(v_i, \theta_i) = w(v_i, \theta_i) - w_\theta(v_i, \theta_i)\psi(\theta_i), \qquad \psi(\theta_i) = \frac{1 - G(\theta_i)}{g(\theta_i)}.$$

The term φ measures marginal revenue from seller i at state (v_i, θ_i), net of information rents. I assume φ is non-decreasing in v_i, is strictly increasing in θ_i and satisfies $\inf \varphi(v_i, \theta_i) < 0 < \sup \varphi(v_i, \theta_i)$, which rules out degenerate cases. Section 6 presents examples satisfying the assumptions.

Independence of types across sellers prevents full surplus extraction via cross-agent correlation as in Crémer and McLean (1988), while independence between each seller's match value and type prevents surplus extraction via type-signal correlation as in McAfee and Reny (1992).

Information Mechanism. The intermediary influences the consumer's choice by controlling the information provided to the consumer about match values. An *information structure* is a conditional distribution

$$\sigma(\cdot \mid \boldsymbol{v}) \in \Delta(\mathcal{S})$$

which maps realized match values $\boldsymbol{v}$ into a distribution over signals $s \in \mathcal{S}$. Given $\boldsymbol{v}$, the intermediary privately sends a signal to the consumer according to $\sigma(\cdot \mid \boldsymbol{v})$.

The intermediary monetizes its information by coupling information provision to the consumer with transfers from sellers. Let Σ be a set of all information structures. An *information mechanism* specifies, for each profile of reported seller types $\boldsymbol{\theta}'$, (i) $\sigma^{\boldsymbol{\theta}'} \in \Sigma$ and (ii) transfers $\boldsymbol{t}(\boldsymbol{v}, \boldsymbol{\theta}') \in \mathbb{R}^N$, provided that all sellers participate.

The intermediary does not provide information directly to sellers. This is without loss of generality, since any seller-directed information can be incorporated into sellers' reports. See Section 5.1 for details.

In principle, because consumers can access all products regardless of seller participation in the mechanism, the intermediary must specify information mechanisms contingent on sellers' participation decisions. However, the intermediary can always induce the consumers to choose from participating sellers' products. Hence, for a revenue-maximizing intermediary, restricting attention to mechanisms that induce full participation with zero payoffs for non-participating sellers is without loss of generality.

Timing of the Game.

1. The intermediary commits to an information mechanism.
2. Match values $\boldsymbol{v}$ and types $\boldsymbol{\theta}$ are realized; the intermediary observes $\boldsymbol{v}$, and sellers observe their own types.
3. Sellers report their types $\boldsymbol{\theta}'$.
4. The consumer observes a signal drawn according to $\sigma^{\boldsymbol{\theta}'}(\cdot \mid \boldsymbol{v})$ and chooses an option; transfers are collected.

2.1 Recommendation Mechanism

A *recommendation mechanism* is an information mechanism in which the intermediary's signal directly recommends one of the products or the outside option. Formally, it is a pair $(\boldsymbol{r}, \boldsymbol{t})$, where $\boldsymbol{r}(\boldsymbol{v}, \boldsymbol{\theta}) \in \Delta(\{0\} \cup \mathcal{N})$ specifies the distribution over recommendations at state $(\boldsymbol{v}, \boldsymbol{\theta})$ and $\boldsymbol{t}(\boldsymbol{v}, \boldsymbol{\theta}) \in \mathbb{R}^N$ specifies transfers from sellers. The component $r_i(\boldsymbol{v}, \boldsymbol{\theta})$ denotes the probability with which option i is recommended at state $(\boldsymbol{v}, \boldsymbol{\theta})$.

Provided that sellers participate and report their types truthfully, the consumer updates his beliefs upon receiving a recommendation and chooses the option with the highest posterior expected match value.

A recommendation mechanism is *obedient* if the consumer always prefers following the recommendation over choosing any alternatives. Formally, obedience requires that for every $i, j \in \{0\} \cup \mathcal{N}$,

$$OB_{j|i} : \int_{\boldsymbol{\mathcal{V}} \times \boldsymbol{\Theta}} v_i r_i(\boldsymbol{v}, \boldsymbol{\theta}) \boldsymbol{F}(d\boldsymbol{v}) \boldsymbol{G}(d\boldsymbol{\theta}) \geq \int_{\boldsymbol{\mathcal{V}} \times \boldsymbol{\Theta}} v_j r_i(\boldsymbol{v}, \boldsymbol{\theta}) \boldsymbol{F}(d\boldsymbol{v}) \boldsymbol{G}(d\boldsymbol{\theta}). \tag{1}$$

The constraint $OB_{j|i}$ is called an *obedience constraint from i to j*: it ensures that, conditional on receiving recommendation i, the consumer weakly prefers choosing i over deviating to j.

Provided that the consumer follows the recommendations, seller i's expected payoff of reporting θ_i' when her type is θ_i is

$$\Pi_i(\theta_i'; \theta_i) = \int_{\boldsymbol{\mathcal{V}} \times \boldsymbol{\Theta}_{-i}} \Big(w(v_i, \theta_i) r_i(\boldsymbol{v}, \theta_i', \boldsymbol{\theta}_{-i}) - t_i(\boldsymbol{v}, \theta_i', \boldsymbol{\theta}_{-i}) \Big) \boldsymbol{F}(d\boldsymbol{v}) \boldsymbol{G}_{-i}(d\boldsymbol{\theta}_{-i}),$$

and that from reporting truthfully as θ_i is

$$\Pi_i^*(\theta_i) = \Pi_i(\theta_i; \theta_i).$$

A recommendation mechanism is individually rational if all sellers participate, that is, for all $i \in \mathcal{N}$ and $\theta_i \in \Theta$,

$$\Pi_i^*(\theta_i) \geq 0, \tag{2}$$

and incentive compatible if all sellers report their types truthfully, that is, for all $i \in \mathcal{N}$ and $\theta_i, \theta_i' \in \Theta$,

$$\Pi_i^*(\theta_i) \geq \Pi(\theta_i'; \theta_i). \tag{3}$$

A recommendation mechanism is *feasible* if it is obedient, individually rational and incentive compatible.

By standard revelation principle arguments from mechanism design (Myerson, 1981) and information design (Bergemann and Morris, 2019), it is without loss of generality to restrict our attention to feasible recommendation mechanisms: Any *outcome* achievable by an information mechanism, namely, the joint distribution over match values, seller transfers, and consumer choices induced in equilibrium, can be replicated by a recommendation mechanism that is obedient, incentive compatible, and individually rational, and conversely. The argument is standard and omitted.

Lemma 1. *An outcome can be attained as a Bayes-Nash equilibrium of an information mechanism if and only if it can be attained by a feasible recommendation mechanism.*

Using a feasible recommendation mechanism, the intermediary maximizes its expected revenue from sellers

$$\int_{\boldsymbol{\mathcal{V}}\times\boldsymbol{\Theta}} \sum_{i\in\mathcal{N}} t_i(\boldsymbol{v},\boldsymbol{\theta})\boldsymbol{F}(d\boldsymbol{v})\boldsymbol{G}(d\boldsymbol{\theta}).$$

To extract revenue, the intermediary must discipline sellers. Unlike in standard mechanism design, this discipline cannot be achieved through direct control of allocations, since the intermediary does not choose the consumer's action. Instead, discipline operates through persuasion: informative recommendations shape the consumer's beliefs and thereby affect choice probabilities. These induced choice probabilities discipline sellers' participation and reporting incentives. When persuasion fails, seller discipline breaks down and revenue extraction collapses.

To illustrate, consider a second-price auction with random tie-breaking. Because sellers do not observe match values, bids depend only on types, so the recommendation conveys no information and is ignored by the consumer. Since bids affect payments but not the choices, bidding zero weakly dominates, and revenue is zero.

Proposition 1. *Under random tie-breaking, second-price auction raises zero revenue.*

2.2 Discussion of the Model

The model captures a form of informational intermediation. The intermediary aggregates information elicited from sellers with its own data and strategically transmits information to the consumer, without directly setting prices or controlling allocations. The intermediary cannot restrict the consumer's access to products, the consumer faces no search costs and has no alternative sources of information, and neither signals nor consumer choices are contractible. Seller pricing is taken as fixed and subsumed into match values and seller types.

These abstractions isolate the role of information provision and persuasion as the intermediary's sole instruments for influencing consumer choice. They are particularly salient in real-time bidding environments in digital advertising, where auctions for individual consumers are run in milliseconds and prices are either absent (as in sponsored links on search engines) or not adjustable based on auction outcomes (as in sponsored products on e-commerce platforms), while consumers retain access to all products through organic search. While incorporating pricing, search, or gatekeeping would enrich the environment, these extensions are not expected to alter the qualitative structure of the optimal information mechanism or the paper's main economic insights.

3 Optimal Recommendation Mechanism

In this section, I characterize an optimal recommendation mechanism.

The consumer's obedience depends only on the informativeness of recommendations, and not on transfers. Thus, given an obedient recommendation rule, transfers can be chosen independently to satisfy incentive compatibility and individual rationality. The standard envelope argument (Myerson (1981)) then implies that the intermediary's problem reduces to maximizing expected virtual willingness to pay.

Lemma 2. *Suppose that a recommendation rule $\boldsymbol{r}$ maximizes*

$$\int_{\boldsymbol{\mathcal{V}}\times\boldsymbol{\Theta}} \sum_{i\in\mathcal{N}} \varphi(v_i,\theta_i) r_i(\boldsymbol{v},\boldsymbol{\theta})\boldsymbol{F}(d\boldsymbol{v})\boldsymbol{G}(d\boldsymbol{\theta}) \tag{4}$$

subject to the obedience constraints (1) and the monotonicity constraints,

$$R_i(\theta_i) = \int_{\boldsymbol{\mathcal{V}}\times\boldsymbol{\Theta}_{-i}} r_i(\boldsymbol{v},\boldsymbol{\theta})\boldsymbol{F}(d\boldsymbol{v})\boldsymbol{G}_{-i}(d\boldsymbol{\theta}_{-i}) \text{ is non-decreasing in } \theta_i \text{ for all } \theta_i\in\Theta,\ i\in\mathcal{N}. \tag{5}$$

Suppose also that, for each $i\in\mathcal{N}$,

$$t_i(\boldsymbol{v},\boldsymbol{\theta}) = w(v_i,\theta_i)r_i(\boldsymbol{v},\boldsymbol{\theta}) - \int_{\underline{\theta}}^{\theta_i} w_\theta(v_i,\tilde{\theta}_i)r_i(\boldsymbol{v},\tilde{\theta}_i,\boldsymbol{\theta}_{-i})d\tilde{\theta}_i. \tag{6}$$

Then, $(\boldsymbol{r},\boldsymbol{t})$ is an optimal recommendation mechanism.

Ignoring the monotonicity constraints and transfers, Lemma 2 recasts the intermediary's mechanism design problem as a Bayesian persuasion problem in which the intermediary maximizes expected virtual willingness to pay subject to obedience. The state dependence of payoffs, multiple consumer actions, and the multi-dimensional state space place the problem outside standard concavification, convex function and duality frameworks. I therefore characterize the solution directly by analyzing the obedience constraints and identifying a monotonicity property that ensure obedience.

Absent obedience constraints, the intermediary directly *allocates* the consumer to the option with the highest virtual willingness to pay. An *optimal allocation rule* is

$$\rho^*(\boldsymbol{v},\boldsymbol{\theta}) = 1 \text{ if } \varphi(v_i,\theta_i) > \max_{j\in\mathcal{N}\setminus\{i\}}(\varphi(v_j,\theta_j),0).$$

Let $\hat{v}_{k|i}(\boldsymbol{r}) = \frac{\int_{\boldsymbol{\mathcal{V}}\times\boldsymbol{\Theta}} v_k r_i(\boldsymbol{v},\boldsymbol{\theta})\boldsymbol{F}(d\boldsymbol{v})\boldsymbol{G}(d\boldsymbol{\theta})}{\int_{\boldsymbol{\mathcal{V}}\times\boldsymbol{\Theta}} r_i(\boldsymbol{v},\boldsymbol{\theta})\boldsymbol{F}(d\boldsymbol{v})\boldsymbol{G}(d\boldsymbol{\theta})}$ be the posterior mean of option k conditioning on the recommendation i under $\boldsymbol{r}$, for i with positive probability, and

$$\bar{v}^* = \hat{v}_{i|i}(\boldsymbol{\rho}^*), \qquad \underline{v}^* = \hat{v}_{i|0}(\boldsymbol{\rho}^*)$$

be the posterior match value of product when it is allocated and when the outside option is allocated under $\boldsymbol{\rho}^*$, respectively. By symmetry, these values are identical across products. Since ρ_i is non-decreasing in and ρ_0 is non-increasing in in v_i,

$$\underline{v}^* \leq \mathbb{E}\left(v_i\right) \leq \overline{v}^*.$$

With obedience constraints, allocations must be implemented through recommendations that the consumer voluntarily follows. Consequently, the intermediary cannot freely distort recommendations to extract revenue: recommended options must remain sufficiently attractive for the consumer to follow them. The optimal rule therefore balances revenue extraction with maintaining credible recommendations.

The next theorem characterizes the optimal recommendation rule and shows how obedience adjusts the recommendations relative to the optimal allocation rule.

Theorem 1. *Let $\boldsymbol{r}^*$ be a recommendation rule such that for each $i \in \mathcal{N}$,*

$$r_i^*(\boldsymbol{v},\boldsymbol{\theta}) = 1 \textit{ if } \varphi(v_i,\theta_i) + \ell_i^*(\boldsymbol{v}) > \max_{j\in\mathcal{N}\setminus\{i\}}(\varphi(v_j,\theta_j) + \ell_j^*(\boldsymbol{v}),0) \tag{7}$$

where

$$\ell_i^*(\boldsymbol{v}) = \begin{cases} 0 & \text{if } v_0 \in [\underline{v}^*, \overline{v}^*] \\ \lambda_0(v_i - v_0) & \text{if } v_0 > \overline{v}^* \\ \lambda_1\left(\frac{1}{N}\sum_{k\in\mathcal{N}} v_k - v_0\right) & \text{if } v_0 < \underline{v}^* \end{cases} \tag{8}$$

where $\lambda_0, \lambda_1 \geq 0$ *are the Lagrange multipliers on the obedience constraints* $OB_{0|i}$ *and* $OB_{i|0}$*, respectively, and are strictly positive whenever the corresponding constraints bind. Then,* $\boldsymbol{r}^*$ *maximizes* (4) *subject to* (1) *and* (5)*.*

By Lemma 2, transfers defined in (6) complete the construction of an optimal recommendation mechanism.

When $\underline{v}^* \leq v_0 \leq \overline{v}^*$, the consumer is nearly ex-ante indifferent across options. The information conveyed by $\boldsymbol{\rho}^*$ suffices for obedience. The optimal recommendation rule coincides with the allocation rule, $\boldsymbol{r}^* = \boldsymbol{\rho}^*$.

When $v_0 > \bar{v}^*$, product recommendations under $\boldsymbol{\rho}^*$ are not sufficiently informative to induce produce choice, so the consumer selects the outside option. This eliminates the intermediary's ability to divert consumer choice towards products and extract revenue. To restore obedience, the optimal rule strengthens the informativeness of product recommendations by tilting them toward states in which match values are high relative to the outside option, generating the adjustment $\lambda_0(v_i - v_0)$, where $\lambda_0 > 0$ is the Lagrange multiplier on the obedience constraint $OB_{0|i}$.

When $v_0 < \bar{v}^*$, outside option recommendations under $\boldsymbol{\rho}^*$ are not sufficiently informative to induce obedience, so the consumer selects a product. This eliminates the intermediary's ability to divert consumer choice away from sellers, (equivalently, to enforce a reserve price). To restore obedience, the optimal rule strengthens the informativeness of outside option recommendations by tying them to states in which aggregate match quality is sufficiently low relative to v_0, generating the adjustment $\lambda_1\left(\frac{1}{N}\sum_{k\in\mathcal{N}} v_k - v_0\right)$, where $\lambda_1 > 0$ is the Lagrange multiplier on $OB_{i|0}$.

It remains to verify obedience across products. Upon recommending product i, the consumer must prefer i to any other products j, $\hat{v}_{i|i}(\boldsymbol{r}) \geq \hat{v}_{j|i}(\boldsymbol{r})$. I provide a sufficient condition on recommendation rules and show the optimal rule satisfies it.

Definition 1. A recommendation rule $\boldsymbol{r}$ is *swap-monotone* in match values if for any

$i, j \in \mathcal{N}$, $\boldsymbol{v}_{-ij} \in \boldsymbol{\mathcal{V}}_{-ij}$ and $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ and $v > v'$,

$$r_i(v_i = v, v_j = v', \boldsymbol{v}_{-ij}, \boldsymbol{\theta}) \geq r_i(v_i = v', v_j = v, \boldsymbol{v}_{-ij}, \boldsymbol{\theta}).$$

Holding all other components fixed, swapping the match values of i and j weakly increases the recommendation probability of the product whose value becomes higher, and decreases that of the lower.

Swap-monotonicity[2] is weaker than coordinatewise monotonicity, which requires r_i to be non-decreasing in v_i and non-increasing in each v_j. Coordinatewise monotonicity implies the obedience by comparing posteriors to priors,

$$\hat{v}_{i|i}(\boldsymbol{r}^*) \geq \mathbb{E}\left(v_i\right) = \mathbb{E}\left(v_j\right) \geq \hat{v}_{j|i}(\boldsymbol{r}^*),$$

which is stronger than necessary. Swap-monotonicity restricts behavior only under value swaps. Since r_i is non-decreasing when values are swapped in favor of i, recommendation of i is more likely in states where $v_i \geq v_j$, yielding the posterior ordering.

Lemma 3. *Any recommendation rule $\boldsymbol{r}$ that is swap-monotone almost surely satisfies all obedience constraints between products.*[3]

The optimal rule $\boldsymbol{r}^*$ selects the product maximizing $\varphi(v_i, \theta_i) + \ell_i^*(\boldsymbol{v})$. Since φ is non-decreasing in v_i and ℓ_i^* depends symmetrically on match values, swapping v_i and v_j in favor of i weakly increases i's adjusted virtual willingness to pay while dropping j's. Thus, $\boldsymbol{r}^*$ is swap-monotone, though it does not satisfy coordinatewise monotonicity, and satisfies obedience across products by Lemma 3. Together with the earlier verifications, Theorem 1 follows.

4 Implementation

This section shows that the optimal recommendation rule can be implemented via a generalization of the second-price auction, and studies its robustness.

[2]Swap-monotonicity and cyclic monotonicity are not comparable in general: cyclic monotonicity is a global convexity condition over arbitrary pairs of states, whereas swap-monotonicity restricts only comparisons induced by exchanging two coordinates. When 2-cycle monotonicity is restricted to such swap pairs, however, swap-monotonicity implies the corresponding inequality.

[3]A partial converse also holds within symmetric recommendation rules: if a symmetric rule satisfies all obedience constraints between products for every F and G, then it must be swap-monotone. The proof uses a two-point perturbation argument and is omitted.

The optimal recommendation rule admits a *scoring representation*. Let $w_i = w(v_i, \theta_i)$. Conditional on v_i, the distribution of w_i is $H(w_i \mid v_i) = G(\tilde{\theta}(v_i, w_i))$, where $\tilde{\theta}(v_i, w_i)$ solves $w(v_i, \tilde{\theta}) = w_i$, with density $h(w_i \mid v_i)$. Define seller i's *score* as

$$\hat{s}_i^*(\boldsymbol{v}, w_i) = w_i - \frac{1 - H(w_i \mid v_i)}{h(w_i \mid v_i)} + \ell_i(\boldsymbol{v}).$$

By change-of-variables, evaluating at $w_i = w(v_i, \theta_i)$ gives $\hat{s}_i^*(\boldsymbol{v}, w(v_i, \theta_i)) = \varphi(v_i, \theta_i) + \ell_i(\boldsymbol{v})$.[4] The optimal rule selects the product with the highest non-negative score:

$$r_i^*(\boldsymbol{v}, \boldsymbol{\theta}) = 1 \text{ if } \hat{s}_i^*(\boldsymbol{v}, w(v_i, \theta_i)) > \max_{j \in \mathcal{N} \setminus \{i\}} \left(\hat{s}_j^*(\boldsymbol{v}, w(v_j, \theta_j)), 0 \right).$$

If sellers observed w_i, the intermediary could elicit it directly and apply the scoring rule. Since they do not, the intermediary must infer w_i from bids and match data.

In a *weighted second maximum willingness to pay auction (wSMWA)*, each seller pays an entry fee τ and bids b_i. The intermediary then infers willingness to pay,

$$\hat{w}_i(\boldsymbol{v}, b_i) = \mathbb{E}\left(w(v_i, \theta_i) \mid \boldsymbol{v}, \bar{w}(\theta_i) = b_i \right),$$

that is, the posterior expectation of seller i's realized valuation given its information $\boldsymbol{v}$, interpreting the bid as reporting the seller's maximum willingness to pay,

$$\bar{w}(\theta_i) = \sup_{v_i \in \mathcal{V}} w(v_i, \theta_i).$$

If sellers bid $b_i = \bar{w}(\theta_i)$, then the bid identifies θ_i, the posterior becomes degenerate, and $\hat{w}_i(\boldsymbol{v}, b_i) = w(v_i, \theta_i)$.

The intermediary assigns score $\hat{s}_i^*(\boldsymbol{v}, \hat{w}_i(\boldsymbol{v}, b_i))$ to each seller, selects the seller with the highest non-negative score (with ties broken by a fixed rule), and charges the smallest willingness to pay at which she would still be selected.

Formally, fix seller i, match values $\boldsymbol{v}$ and opponents' bids $\boldsymbol{b}_{-i}$. Let $\hat{\boldsymbol{w}}_{-i} =$

[4]Because $w(v, \theta)$ is strictly increasing in θ, the change-of-variables formula implies

$$h(w \mid v) = \frac{g(\tilde{\theta}(v, w))}{w_\theta(\tilde{\theta}(v, w), v)}.$$

Substituting into $w - (1 - H)/h$ yields φ under the mapping $w = w(v, \theta)$.

$\hat{\boldsymbol{w}}_{-i}(\boldsymbol{v}, \boldsymbol{b}_{-i})$. Define

$$\hat{p}_i^*(\boldsymbol{v}, \hat{\boldsymbol{w}}_{-i}) = \inf\left\{x \in \mathcal{W} \mid \hat{s}_i^*(\boldsymbol{v}, x, \hat{\boldsymbol{w}}_{-i}) \geq \max\left\{0, \max_{j \neq i} \hat{s}_j^*(\boldsymbol{v}, x, \hat{\boldsymbol{w}}_{-i})\right\}\right\}. \quad (9)$$

with $\hat{p}_i(\boldsymbol{v}, \hat{\boldsymbol{w}}_{-i}) = \sup \mathcal{W}$ if the set is empty. Seller i wins if and only if

$$\hat{w}_i(\boldsymbol{v}, b_i) > \hat{p}_i^*(\boldsymbol{v}, \hat{\boldsymbol{w}}_{-i}),$$

and wins with a fixed tie-breaking probability when equality holds. Conditional on winning, she pays $\hat{p}_i(\boldsymbol{v}, \hat{\boldsymbol{w}}_{-i})$; otherwise, she pays zero.

Because selection depends on both types and match values, even the lowest type $\underline{\theta}$ may win with positive probability and thus earn a positive expected payoff absent an entry fee. Define

$$p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) = \hat{p}_i^*(\boldsymbol{v}, \boldsymbol{w}_{-i}(\boldsymbol{v}_{-i}, \boldsymbol{\theta}_{-i}))$$

as the payment conditional on winning when all other sellers' bids reveal their types. Seller $\underline{\theta}$ wins if and only if

$$w(v_i, \underline{\theta}) > p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})$$

in which case she pays $p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})$. Her expected payoff absent an entry fee is therefore

$$\tau^* = \int_{\boldsymbol{\mathcal{V}} \times \boldsymbol{\Theta}_{-i}} (w(v_i, \underline{\theta}) - p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})) \cdot \mathbf{1}\{w(v_i, \underline{\theta}) > p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})\} \boldsymbol{F}(d\boldsymbol{v}) \boldsymbol{G}_{-i}(d\boldsymbol{\theta}_{-i}),$$

and the intermediary extracts this surplus by charging an entry fee τ^*.

Theorem 2. *A weighted second maximum willingness to pay auction with an entry fee τ^* implements an optimal recommendation mechanism, under which each seller bids $b_i^*(\theta_i) = \bar{w}(\theta_i)$ weakly dominantly.*

In equilibrium, each seller bids her maximum willingness to pay $\bar{w}(\theta_i)$, which is strictly increasing in type and therefore fully revealing. Sellers are willing to reveal this information because, conditional on inferred valuations, recommendation and payments follow a weakly dominant, second-price logic.

Given revealing bids, the intermediary combines sellers' reports with its own information to infer willingness to pay and determine the winner and payment. The wSMWA can thus be interpreted as a generalized second-price auction in which willingness to pay is not directly observed by sellers or the intermediary. Unlike a standard second-price auction, where bids directly represent valuations, valuations here are reconstructed by aggregating sellers' and the intermediary's information, and second-price logic is applied in this inferred valuation space.

Interpreting bids as sellers' maximum willingness to pay, as in the wSMWA, is only one way of eliciting private information. More generally, the intermediary may use any bijection $\tilde{\theta} : \mathbb{R} \to \mathbb{R}$ to interpret bids. Under this procedure, a bid b_i is treated as reporting type $\tilde{\theta}(b_i)$. An inferred willingness to pay is

$$\tilde{w}_i(\boldsymbol{v}, b_i) = \mathbb{E}\left(w(v_i, \theta_i) \mid \boldsymbol{v}, \tilde{\theta}^{-1}(\theta_i) = b_i\right).$$

The intermediary then applies the same scoring rule $\hat{\boldsymbol{s}}^*$, payment rule $\hat{\boldsymbol{p}}^*$ and entry fee τ^* as in wSMWA, using the inferred willingness to pay $\tilde{w}_i(\boldsymbol{v}, b_i)$. We refer to any auction constructed in this way as a *$\tilde{\theta}$-auction*.

In this $\tilde{\theta}$-auction, $\tilde{b}(\theta) = \tilde{\theta}^{-1}(\theta)$ is weakly dominant by the same second-price logic as in Theorem 2. Under this strategy, inferred willingness to pay equals the true value

$$\tilde{w}_i(v, \tilde{b}_i(\theta_i)) = w(v_i, \theta_i).$$

Outcomes coincide with wSMWA: i wins if $w(v_i, \theta_i) > p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})$ and pays $p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})$.

The wSMWA is a $\tilde{\theta}$-auction in which $\tilde{\theta}$ maps bids to maximum willingness to pay. The next proposition shows that wSMWA is the unique $\tilde{\theta}$-auction satisfying two natural properties of equilibrium bids and payments conditional on winning.

Proposition 2. *Fix a scoring rule $\hat{\boldsymbol{s}}$, payment rule $\hat{\boldsymbol{p}}^*$ and entry fee τ^* as defined above. For each $i \in \mathcal{N}$ and $\theta_i \in \Theta$, let*

$$\mathcal{C}(\theta_i) = \{(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \mid w(v_i, \theta_i) > p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})\},$$

be the states in which type θ_i wins. The wSMWA is the unique auction among all $\tilde{\theta}$-auctions satisfying the following: for every $i \in \mathcal{N}$ and $\theta_i \in \Theta$ such that $\mathcal{C}(\theta_i) \neq \emptyset$,

Maximum Bid Property: $\tilde{b}(\theta_i) \geq p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})$ *for all* $(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \in \mathcal{C}(\theta_i)$.

Tight Bound Property: $\tilde{b}(\theta_i) = p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})$ *for some* $(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \in \mathcal{C}(\theta_i)$.

The Maximum Bid Property guarantees that a seller is never charged more than her equilibrium bid. The Tight Bound Property ensures that this upper bound is meaningful: there exists a state in which the seller pays her full equilibrium bid. Together, these two properties characterize auctions in which bids act as tight upper bounds on equilibrium payments, conditional on winning. Proposition 2 shows that this property uniquely pins down the wSMWA among all $\tilde{\theta}$-auctions. Since both properties hold in a standard second-price auction, wSMWA can be viewed as its unique extension to environments with partial information.

Modern digital advertising platforms ask advertisers to submit a maximum willingness to pay bid,[5] representing the highest amount they are willing to pay and can be charged, rather than a state-contingent valuation. The platform then combines the bid with proprietary relevance or match-quality signals to determine recommendation and payment. The wSMWA provides a mechanism design foundation for this institutional structure. When sellers are only partially informed about their and the platform holds additional match information, requesting a maximum willingness to pay is not ad hoc: it emerges as the unique dominant-strategy implementation under tight upper bounds on payments.

4.1 Robustness

This subsection examines the robustness of wSMWA across different environments and compares it to other auction formats.

The consumer's choice is *allocative* if the intermediary directly assigns the consumer to sellers, so obedience constraints are irrelevant; it is *persuasive* if the consumer instead selects the option that maximizes his posterior expected value given the recommendation. Sellers are *fully informed* if they observe their realized willingness to pay, either by observing match values or because match values are do not affect willingness to pay, and *partially informed* otherwise. The intermediary is *informed* if it observes match values, and *uninformed* otherwise.

I restrict attention to environments in which match values are economically relevant: either the intermediary is informed, or sellers are fully informed while match values affect willingness to pay. This rules out the environment with partially informed sellers and uninformed intermediary. I refer to this condition as *Relevance*.

[5] See Google Ads Help, *How the Google Ads auction works*, https://support.google.com/google-ads/answer/6167118, explaining that advertisers submit a maximum bid representing the highest amount they are willing to pay, which is combined with proprietary quality and relevance signals to determine both ad ranking and payment. Platforms use various terms for the bid, including "maximum bid," "bid cap," or "maximum cost per result."

An *environment* is defined as a specification of consumer behavior, seller information, and intermediary information. The *baseline* environment features persuasive consumer choice, partially informed sellers, and an informed intermediary. At the opposite extreme, an *auction-like* environment features allocative consumer choice, fully informed sellers and uninformed intermediary.

To compare across environments, I impose the following regularity assumption.

Assumption 1. *Let $w = w(v, \theta)$ be a random variable with cumulative distribution function H and density h. Assume that the unconditional virtual willingness to pay*

$$\varphi^U(w) = w - \frac{1 - H(w)}{h(w)}$$

and the posterior mean

$$\hat{v}(w) = \mathbb{E}\left(v \mid w\right)$$

are strictly increasing in true willingness to pay w.

The *second-price auction (SPA)* selects the highest bidder and charges the second-highest bid (or a reserve). The *weighted second-price auction (wSPA)*, introduced by Gomes (2014), treats bids as willingness to pay and applies the optimal scoring rule. The wSMWA instead infers willingness to pay from bids and available information, interpreting bids as maximum willingness to pay given private information, and applies the same rule to the inferred values. See Appendix 3 for the formal construction.

An auction format is *optimal* if it implements the optimal recommendation mechanism. It *reduces to* another format if their recommendation and payment rules coincide for every realization of the intermediary's information and bids.

Theorem 3. *Under Assumption 1 and Relevance:*

1. *SPA is optimal if and only if the environment is auction-like.*
2. *wSPA is optimal if and only if sellers are fully informed.*
3. *wSMWA is optimal in all environments.*
4. *wSMWA reduces to wSPA if and only if sellers are fully informed.*
5. *wSPA and wSMWA reduce to SPA if and only if the environment is auction-like.*

The classification highlights the distinct sources of failure of standard auction formats. The SPA is optimal in auction-like environments, where the problem reduces to Myerson (1981) setting in which bids fully capture willingness to pay and no informational adjustment is required. Relaxing any of these conditions renders SPA suboptimal. With persuasive consumer choice, bid-based recommendations are insufficient to induce obedience. With partially informed sellers, bids no longer reflect true willingness to pay. With an informed intermediary, restricting recommendations to bids prevents the mechanism from exploiting match value information.

The wSPA extends SPA by incorporating intermediary information into the scoring rule, but it continues to interpret bids as realized willingness to pay. When sellers are fully informed, this interpretation is correct, and wSPA is optimal regardless of consumer choice and intermediary information. Once sellers are partially informed, however, bids reveal only private types rather than realized valuations, and wSPA fails to recover true willingness to pay.

The wSMWA extends wSPA by inferring willingness to pay from bids and available information, instead of presuming bids reveal it. By correctly reconstructing realized willingness to pay, it is optimal in all environments.

When sellers are fully informed, bids directly equal realized willingness to pay and wSMWA reduces to wSPA. In auction-like environments, where neither persuasion nor additional information is present, both auction rules reduce to the SPA.

5 Information Access

This section analyzes how changes in information access affect equilibrium outcomes under the optimal recommendation mechanism. I consider three ways in which access to information may change: direct sales or disclosure to sellers, sellers' information acquisition, and regulatory restrictions on data access.

5.1 Direct Sales and Disclosure

The baseline model studies indirect monetization of data through recommendations. A natural question is whether the intermediary can raise additional revenue by selling this information directly to sellers before operating the recommendation mechanism.

Suppose the intermediary may sell match value information before operating the recommendation mechanism. It offers a menu of information structures to sellers and then proceeds with recommendations and transfers.

Any equilibrium outcome of this procedure induces a joint distribution over match values, recommendations, and total payments (including upfront information charges). By the revelation principle, this outcome can be implemented by a recommendation mechanism. Since the optimal recommendation mechanism already maximizes revenue, selling match information separately cannot strictly increase revenue.

Disclosure is a special case in which the price of information is zero, so the same argument applies when the intermediary freely discloses match information to sellers. This establishes the following result.

Proposition 3. *Allowing the intermediary to sell or disclose match information to sellers does not change the optimal recommendation mechanism.*

The irrelevance of disclosure contrasts with environments in which the intermediary does not observe valuation-relevant information or faces constraints on mechanism design, where information release can strictly increase revenue (e.g., Eső and Szentes, 2007; Bergemann, Heumann, and Morris, 2021). In the present setting, the intermediary already conditions recommendations and transfers on privately observed match values, so disclosure does not expand the set of implementable outcomes.

Consequently, regulations that restrict the sale of match data to sellers, or limits or mandates disclosure to them, do not affect equilibrium recommendations, payments, or welfare under the optimal mechanism.

5.2 Seller Information Acquisition

Allowing sellers to acquire signals about match values does not change the optimal recommendation mechanism.

Suppose that before reporting their types, sellers observe additional private signals correlated with match values, for example through data collection technologies or third-party data brokers. The intermediary commits to a recommendation mechanism that may condition on match values and reported seller information.

Although expanding sellers' information sets can tighten incentive constraints, seller signals do not enter the intermediary's objective or the consumer's obedience constraints conditional on match values. Because match values are already observed and determine the optimal allocation rule, additional seller information does not expand the set of revenue-relevant outcomes.

Consequently, costly acquisition of such information cannot be profitable in equilibrium, and restrictions on sellers' ability to collect or purchase match-relevant data, such as limits on tracking technologies or third-party data markets, are non-binding. When the intermediary already implements optimal targeting through its allocation rule, additional seller-side targeting has no effect on market outcomes.

5.3 Consumer Welfare and Match Data Regulation

I analyze consumer welfare and restrictions on match data access.

Absent recommendations, the consumer selects a *default option* that maximizes his prior expected match value. Consumer welfare, the black curve in Figure 1, is,

$$CW^D = \max(\mathbb{E}(v_i), v_0).$$

Under an obedient recommendation rule $\boldsymbol{r}$, consumer welfare equals the expected match value of recommended options,

$$CW(\boldsymbol{r}) = \int_{\boldsymbol{\mathcal{V}}\times\boldsymbol{\Theta}} \sum_{i\in\{0\}\cup\mathcal{N}} v_i r_i(\boldsymbol{v},\boldsymbol{\theta})\boldsymbol{F}(d\boldsymbol{v})\boldsymbol{G}(d\boldsymbol{\theta}).$$

Access to match values allows recommendations to be informative about fit relative to the outside option and across ex-ante symmetric products. When recommendations are sufficiently informative, the consumer strictly prefers to follow them and welfare strictly exceeds CW^D.

The welfare effect depends on v_0. If $\underline{v}^* < v_0 < \overline{v}^*$, ex-ante differences among options are small, any recommendation is sufficiently informative, $CW(\boldsymbol{r}^*) > CW^D$. If $v_0 \geq \overline{v}^*$, product recommendations become just sufficiently informative to match the outside option, yielding $CW(\boldsymbol{r}^*) = v_0 = CW^D$. If $v_0 \leq \underline{v}^*$, product recommendations are strictly informative across ex-ante symmetric products, so $CW(\boldsymbol{r}^*) > CW^D$.

Consumer welfare is depicted in Figure 1 by the blue curve.

Theorem 4. *The optimal recommendation mechanism increases (does not change) consumer welfare if and only if $v_0 < \overline{v}^*$ ($v_0 \geq \overline{v}^*$).*

Theorem 4 admits a natural interpretation in terms of match-data regulation. Prohibiting the use of consumer–product match data prevents conditioning on fit, rendering recommendations uninformative and reducing welfare to the no-recommendation

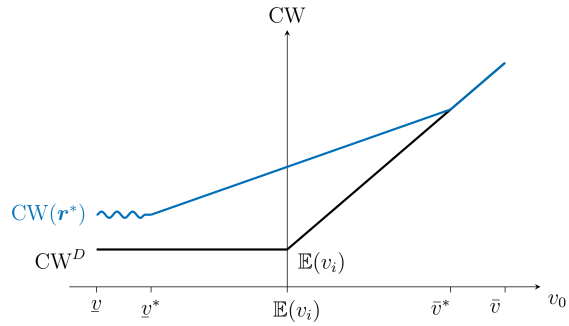

Figure 1: Comparison of consumer welfare with and without the optimal recommendation rule.

benchmark. Such restrictions cannot increase consumer welfare and strictly reduce it for consumers with poor outside options.

5.4 Seller Data Regulation

This subsection studies whether restricting the intermediary's access to seller data improves consumer welfare by comparing the baseline environment with one in which the intermediary observes seller data revealing sellers' private information.

To facilitate the analysis, I impose an additional set of assumptions.

Assumption 2. *F is absolutely continuous on $[\underline{v}, \overline{v}]$, and G is strictly log-concave. The functions w and φ are C^2 with strictly positive partial derivatives. For $\phi \in \{w, \varphi\}$, let $m^\phi(v,\theta) = \frac{\phi_v(v,\theta)}{\phi_\theta(v,\theta)}$, and*

$$m^w(v,\theta) \geq m^\varphi(v,\theta), \qquad m^\varphi_\theta(v,\theta) \leq 0.$$

The ratio m^φ captures the relative weight placed on v versus θ under φ. The inequality $m^w \geq m^\varphi$ states that the virtual willingness to pay places less importance to match value than its true counterpart. The condition $m^\varphi_\theta \leq 0$ implies that the conditional distribution of φ given v satisfies a monotone reverse hazard rate property in v, so higher v shifts the distribution of φ upward in that order.

Without seller data, the intermediary maximizes virtual willingness to pay,

$$\int_{\boldsymbol{\mathcal{V}}\times\boldsymbol{\Theta}} \sum_{i\in\{0\}\cup\mathcal{N}} \varphi(v_i,\theta_i) r_i(\boldsymbol{v},\boldsymbol{\theta})\boldsymbol{F}(d\boldsymbol{v})\boldsymbol{G}(d\boldsymbol{\theta}),$$

subject to obedience.

With seller data, the intermediary observes seller types, eliminates information rents, and maximizes true willingness to pay subject to obedience:

$$\int_{\boldsymbol{\mathcal{V}}\times\boldsymbol{\Theta}} \sum_{i\in\{0\}\cup\mathcal{N}} w(v_i,\theta_i) r_i(\boldsymbol{v},\boldsymbol{\theta})\boldsymbol{F}(d\boldsymbol{v})\boldsymbol{G}(d\boldsymbol{\theta}). \tag{10}$$

Let $\boldsymbol{\rho}^S$ be an optimal allocation rule and $\boldsymbol{r}^S$ be an optimal recommendation rule. Let

$$\overline{v}^S = \hat{v}_{i|i}(\boldsymbol{\rho}^S)$$

be the posterior match value of a recommended product with seller data.

The outside option is recommended only when obedience binds, so welfare is

$$CW(\boldsymbol{r}^S) = \max\{\overline{v}^S, v_0\}.$$

The welfare comparison is governed by $\overline{v}^*$, $\overline{v}^S$, and their interaction with v_0.

Theorem 5. *Suppose Assumption 2 holds.*

1. *If $\bar{v}^S \geq \bar{v}^*$, then*

$$CW(\boldsymbol{r}^S) \geq CW(\boldsymbol{r}^*) \quad \textit{for all } v_0,$$

with strict inequality for all $v_0 < \overline{v}^S$.

2. *If $\bar{v}^S < \bar{v}^*$, then there is $\underline{v}^S \in (\underline{v}^*, \overline{v}^*)$ such that*

$$\begin{aligned} CW(\boldsymbol{r}^S) &> CW(\boldsymbol{r}^*) \quad \textit{for all } v_0 < \underline{v}^S, \\ CW(\boldsymbol{r}^S) &< CW(\boldsymbol{r}^*) \quad \textit{for all } v_0 \in (\underline{v}^S, \overline{v}^*), \\ CW(\boldsymbol{r}^S) &= CW(\boldsymbol{r}^*) \quad \textit{for all } v_0 \geq \overline{v}^*. \end{aligned}$$

For any symmetric recommendation rule $\boldsymbol{r}$,

$$CW(\boldsymbol{r}) = \hat{v}_{i|i}(\boldsymbol{r})(1 - \Pr(r_0)) + v_0 \Pr(r_0), \tag{11}$$

Welfare depends on two sufficient statistics: product-recommendation informativeness $\hat{v}_{i|i}(\boldsymbol{r})$ and the probability of recommending the outside option $\Pr(r_0)$. In particular, $\overline{v}^*$ and $\overline{v}^S$ measure the informativeness in the respective environment.

In the baseline, disciplining sellers requires *diversion*: recommendations are dis-

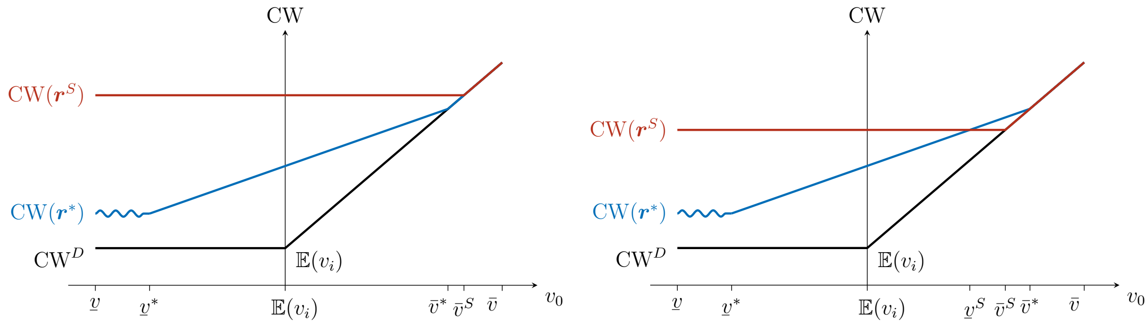

(a) $\overline{v}^S \geq \overline{v}^*$ (b) $\overline{v}^S < \overline{v}^*$

Figure 2: Comparison of consumer welfare in the baseline and seller data environment

torted away from true willingness to pay maximizer to the outside option or rivals.

Seller data eliminates these distortions, and affects welfare through two channels. It reduces outside-option diversion, lowering $\Pr(r_0)$, and alters product-recommendation informativeness: expanding product recommendations into lower match states reduces $\overline{v}^S$, while aligning recommendations with match values instead of information rents increases $\overline{v}^S$. The net effect determines whether $\overline{v}^S \geq \overline{v}^*$ or not.

Welfare is governed by these two channels. When v_0 is low, recommending the outside option sacrifices substantial match value, so eliminating such diversion strictly increases welfare regardless of the informational effect. When v_0 is high, welfare equals v_0 in both environments, so that seller data has no effect. For intermediate v_0, the welfare effect depends on the informational channel—namely, whether seller data increases product recommendation informativeness ($\overline{v}^S \geq \overline{v}^*$).

Restricting seller data access restores information rents and diversion distortions. Consumers with sufficiently low outside options are strictly worse off. If seller data increases product informativeness ($\overline{v}^S \geq \overline{v}^*$), restricting access benefits no consumer. Figure 2 illustrates the welfare comparison in both cases.

6 Examples

This section presents examples satisfying the maintained assumptions. Throughout, v has compact support $\mathcal{V}$, $\theta \in [\underline{\theta}, \overline{\theta}]$, v and θ are independent, and G has monotone hazard rate. Moreover, $\inf_{v_i \in \mathcal{V}, \theta_i \in \Theta} \varphi(v_i, \theta_i) < 0 < \sup_{v_i \in \mathcal{V}, \theta_i \in \Theta} \varphi(v_i, \theta_i)$.

Example 1. Consider $w(v, \theta) = \theta$. Assumptions in Section 2 and Assumption 1 hold trivially. The optimal mechanism and its implementation remain valid. □

Example 2. Consider $w(v,\theta) = v + \theta$, each v and θ drawn from a uniform distribution over compact intervals. Assumptions in Section 2 hold trivially. Since w is a sum of log-concave distributions, w has a log-concave density and satisfies monotone likelihood ratio property in v, implying Assumption 1. Assumption 2 follows from $(\log G)'' < 0$, $m^w = 1$ and $m^\varphi = 0$. □

Example 3. Consider $w(v,\theta) = c(av + b\theta)$, with log-concave densities and compact interval supports, where $a, b > 0$ and c is a C^2 function with $c \geq 0$, $c' > 0$ and $c'' \leq 0$. This satisfies assumptions in Section 2 and Assumption 1. If, in addition, $(\log c')'' > 0$ and $\psi'' \leq 0$, then this satisfies Assumption 2. □

7 Conclusion

This paper studies how an intermediary optimally monetizes data by selling recommendations. Because the intermediary cannot directly control consumer choice, recommendations must remain sufficiently informative for the consumer to follow while disciplining sellers' incentives. I characterize the optimal recommendation mechanism and show that it recommends the option with the highest adjusted virtual willingness to pay, where the adjustment reflects the need to maintain credible recommendations.

The optimal mechanism can be implemented in weakly dominant strategies by a weighted second maximum willingness to pay auction. This auction provides a robust generalization of the second-price auction to environments with persuasion, partial seller information, and intermediary-held data, while uniquely preserving the property that equilibrium payments are tightly bounded above by bids.

More broadly, the analysis provides a mechanism design foundation for data-driven recommendation markets. The framework highlights how persuasion constraints interact with revenue extraction and information asymmetries in digital platforms. Extending the framework to settings with competing intermediaries, dynamic recommendations, or endogenous pricing are natural directions for future research.

A Proofs for Section 3

A.1 Proof for Lemma 3

Fix $i \neq j \in \mathcal{N}$, and fix $\boldsymbol{v}_{-ij}$ and $\boldsymbol{\theta}$. Let

$$I = \int_{\boldsymbol{\mathcal{V}}_{ij}} (v_i - v_j)\, r_i(v_i, v_j, \boldsymbol{v}_{-ij}, \boldsymbol{\theta})\, F(dv_i)F(dv_j).$$

Since v_i and v_j are i.i.d., swapping variables gives $I = -\int_{\boldsymbol{\mathcal{V}}_{ij}} (v_i - v_j)\, r_i(v_j, v_i, \boldsymbol{v}_{-ij}, \boldsymbol{\theta})\, F(dv_i)F(dv_j)$. Hence, for $\Delta = r_i(v_i, v_j, \boldsymbol{v}_{-ij}, \boldsymbol{\theta}) - r_i(v_j, v_i, \boldsymbol{v}_{-ij}, \boldsymbol{\theta})$,

$$2I = \int_{\boldsymbol{\mathcal{V}}_{ij}} (v_i - v_j)\, \Delta\, F(dv_i)F(dv_j).$$

On $\{v_i > v_j\}$, swap-monotonicity implies $\Delta \geq 0$; on $\{v_i < v_j\}$, $\Delta \leq 0$. The integrand is non-negative and $I \geq 0$. Integrating over $(\boldsymbol{v}_{-ij}, \boldsymbol{\theta})$ yields the result.

A.2 Proof for Theorem 1

Since $\boldsymbol{r}^*$ solves the pointwise Lagrangian problem, it suffices to verify monotonicity (5) and obedience. Each r_i^* is non-decreasing in θ_i, implying (5). Because $\boldsymbol{r}^*$ is swap-monotone for all v_0, obedience across products holds by Lemma 3.

If $\underline{v}^* \leq v_0 \leq \overline{v}^*$, then $\boldsymbol{r}^* = \boldsymbol{\rho}^*$. Since r_i^* is non-decreasing and ρ_0^* is non-increasing in v_i, so $\hat{v}_{i|0}(\boldsymbol{\rho}^*) = \underline{v}^* \leq v_0 \leq \overline{v}^* = \hat{v}_{i|i}(\boldsymbol{\rho}^*)$, implying $OB_{i|0}$ and $OB_{0|i}$.

If $v_0 > \overline{v}^*$, then r_0^* is non-increasing in v_i, so $\hat{v}_{i|0}(\boldsymbol{r}^*) \leq \mathbb{E}\,(v_i) \leq \overline{v}^* < v_0$, implying $OB_{i|0}$. Moreover, $OB_{0|i}$ binds: otherwise $\lambda_0 = 0$ and $\boldsymbol{r}^* = \boldsymbol{\rho}^*$, violating obedience.

If $v_0 < \underline{v}^*$, then r_i^* is non-decreasing in v_i, so $\hat{v}_{i|i}(\boldsymbol{r}^*) \geq \mathbb{E}\,(v_i) \geq \underline{v}^* > v_0$, implying $OB_{0|i}$. Moreover, $OB_{i|0}$ binds by the same argument.

B Proofs for Section 4

B.1 Proof for Theorem 2

Fix seller i with type θ_i, $(\boldsymbol{v}, \boldsymbol{b}_{-i})$ and $\hat{\boldsymbol{w}}_{-i} = \hat{\boldsymbol{w}}_{-i}(\boldsymbol{v}, \boldsymbol{b}_{-i})$. By (9), seller i wins if and only if $\hat{w}(\boldsymbol{v}, b_i) > \hat{p}_i^*(\boldsymbol{v}, \hat{\boldsymbol{w}}_{-i})$. Seller i's payoff from bid b_i is $w(v_i, \theta_i) - \hat{p}_i^*(\boldsymbol{v}, \hat{\boldsymbol{w}}_{-i})$ if

$\hat{w}(\boldsymbol{v}, b_i) > \hat{p}_i^*(\boldsymbol{v}, \hat{\boldsymbol{w}}_{-i})$ and 0 otherwise. If she bids $b_i^*(\theta_i) = \bar{w}(\theta_i)$, then $\hat{w}(v_i, b_i^*(\theta_i)) = w(v_i, \theta_i)$, so she wins iff her payoff is positive. Thus, $b_i^*(\theta_i)$ is weakly dominant.

B.2 Proof for Proposition 2

Fix a mapping $\tilde{\theta} : \mathbb{R} \to \mathbb{R}$, and consider the associated $\tilde{\theta}$-auction. Under equilibrium bidding $\tilde{b}(\theta_i) = \tilde{\theta}^{-1}(\theta_i)$, inferred willingness to pay equals $w(v_i, \theta_i)$. Thus outcomes coincide with wSMWA: seller i wins iff $w_i(v_i, \theta_i) > p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})$, paying $p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i})$.

I first characterize the maximal equilibrium payment.

Lemma 4. *For every $i \in \mathcal{N}$ and $\theta_i \in \Theta$ such that $\mathcal{C}(\theta_i) \neq \emptyset$,*

$$\sup_{(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \in \mathcal{C}(\theta_i)} p_i^*(\boldsymbol{v}, \boldsymbol{\theta}) = \bar{w}(\theta_i).$$

Proof. For any $(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \in \mathcal{C}(\theta_i)$, $p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) < w(v_i, \theta_i) \leq \bar{w}(\theta_i)$, hence,

$$\sup_{(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \in \mathcal{C}(\theta_i)} p_i^*(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \leq \bar{w}(\theta_i).$$

To show the reverse inequality, define each seller k's score in $(\boldsymbol{v}, \boldsymbol{\theta})$-space by

$$s_k^*(\boldsymbol{v}, \theta_k) = \hat{s}_k^*(\boldsymbol{v}, w(v_k, \theta_k)).$$

Since $\mathcal{C}(\theta_i) \neq \emptyset$, there exists $(\boldsymbol{v}^*, \boldsymbol{\theta}_{-i}^*)$ with $s_i^*(\boldsymbol{v}^*, \theta_i) > 0$.

Fix $\epsilon > 0$ and consider a state $(\bar{\boldsymbol{v}}, \boldsymbol{\theta}^\epsilon)$ such that $v_k = \overline{v}$ for all $k \in \mathcal{N}$, $\theta_i^\epsilon = \theta_i$ and $\theta_k^\epsilon = \theta_i - \epsilon$ for all $k \neq i$. Since s_i^* is non-decreasing in $\boldsymbol{v}$,

$$s_i^*(\bar{\boldsymbol{v}}, \theta_i) \geq s_i^*(\boldsymbol{v}^*, \theta_i) > 0$$

and since s_k^* is strictly increasing and continuous in type, for all sufficiently small $\epsilon > 0$, $s_k^*(\bar{\boldsymbol{v}}, \theta_i - \epsilon) \in (0, s_i^*(\bar{\boldsymbol{v}}, \theta_i))$ for each $k \neq i$. Hence, seller i wins at $(\bar{\boldsymbol{v}}, \boldsymbol{\theta}^\epsilon)$, and

$$p_i^*(\bar{\boldsymbol{v}}, \boldsymbol{\theta}^\epsilon) = w(\overline{v}, \theta_i - \epsilon) = \bar{w}(\theta_i - \epsilon).$$

Hence, $\sup_{(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \in \mathcal{C}(\theta_i)} p_i^*(\boldsymbol{v}, \boldsymbol{\theta}) \geq p_i^*(\bar{\boldsymbol{v}}, \boldsymbol{\theta}^\epsilon)$. Letting $\epsilon \to 0$ gives the inequality:

$$\sup_{(\boldsymbol{v}, \boldsymbol{\theta}_{-i}) \in \mathcal{C}(\theta_i)} p_i^*(\boldsymbol{v}, \boldsymbol{\theta}) \geq \lim_{\epsilon \to 0} p_i^*(\bar{\boldsymbol{v}}, \boldsymbol{\theta}^\epsilon) = \lim_{\epsilon \to 0} \bar{w}(\theta_i - \epsilon) = \bar{w}(\theta_i).$$

□

Let $i \in \mathcal{N}$ and $\theta_i \in \Theta$ such that $\mathcal{C}(\theta_i) \neq \emptyset$. The Maximum Bid Property implies

$$\tilde{b}(\theta_i) \geq \sup_{(\boldsymbol{v},\boldsymbol{\theta}_{-i})\in\mathcal{C}(\theta_i)} p_i^*(\boldsymbol{v},\boldsymbol{\theta}).$$

By the Tight Bound Property, there is $(\boldsymbol{v},\boldsymbol{\theta}_{-i})$ such that $\tilde{b}(\theta_i) = p(\boldsymbol{v},\boldsymbol{\theta}_{-i})$, implying

$$\tilde{b}(\theta_i) \leq \sup_{(\boldsymbol{v},\boldsymbol{\theta}_{-i})\in\mathcal{C}(\theta_i)} p_i^*(\boldsymbol{v},\boldsymbol{\theta}).$$

Combining the inequalities and Lemma 4, bidding rule is uniquely pinned down to

$$\tilde{b}(\theta_i) = \sup_{(\boldsymbol{v},\boldsymbol{\theta}_{-i})\in\mathcal{C}(\theta_i)} p_i^*(\boldsymbol{v},\boldsymbol{\theta}) = \bar{w}(\theta_i).$$

B.3 Proof for Theorem 3

The proof proceeds in three steps. First, I construct a unified framework that accommodates all environments of interest. Second, I characterize the optimal scoring rule in each environment. Third, I establish a general implementation principle and apply this principle to classify SPA, wSPA, and wSMWA.

Throughout this subsection, Assumption 1 and Relevance are maintained.

Unified Framework. To treat informed and uninformed intermediary environments within a single framework, I introduce an augmented information space.

Let $\boldsymbol{\mathcal{Z}} = \boldsymbol{\mathcal{V}} \cup \{\emptyset\}$. The intermediary observes information about match values through a mapping $\boldsymbol{z} : \boldsymbol{\mathcal{V}} \to \boldsymbol{\mathcal{Z}}$ defined by

$$\boldsymbol{z}(\boldsymbol{v}) = \begin{cases} \boldsymbol{v} & \text{if the intermediary is informed} \\ \emptyset & \text{if the intermediary is uninformed} \end{cases}.$$

Define a forward map $T : \boldsymbol{\mathcal{V}} \times \boldsymbol{\Theta} \to \boldsymbol{\mathcal{Z}} \times \boldsymbol{\mathcal{W}}$ as

$$T(\boldsymbol{v},\boldsymbol{\theta}) = (\boldsymbol{z}(\boldsymbol{v}), \boldsymbol{w}(\boldsymbol{v},\boldsymbol{\theta}))$$

Let $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{t}}) : \boldsymbol{\mathcal{Z}} \times \boldsymbol{\mathcal{W}} \to \Delta(\{0\} \cup \mathcal{N}) \times \mathbb{R}^N$ and $\hat{\boldsymbol{s}} : \boldsymbol{\mathcal{Z}} \times \boldsymbol{\mathcal{W}} \to \mathbb{R}^N$ be a recommendation mechanism and a scoring rule defined on $(\boldsymbol{z}, \boldsymbol{w})$-space. For each recommendation mechanism $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{t}})$ defined on $(\boldsymbol{z}, \boldsymbol{w})$, there is an associated recommendation rule $(\boldsymbol{r}, \boldsymbol{t})$ and scoring rule $\mathbf{s}$ defined on $(\boldsymbol{v}, \boldsymbol{\theta})$-space

$$\boldsymbol{r}(\boldsymbol{v}, \boldsymbol{\theta}) = \hat{\boldsymbol{r}}(T(\boldsymbol{v}, \boldsymbol{\theta})), \qquad \boldsymbol{t}(\boldsymbol{v}, \boldsymbol{\theta}) = \hat{\boldsymbol{t}}(T(\boldsymbol{v}, \boldsymbol{\theta})), \qquad \mathbf{s}(\boldsymbol{v}, \boldsymbol{\theta}) = \hat{\boldsymbol{s}}(T(\boldsymbol{v}, \boldsymbol{\theta}))$$

The above construction extends all objects from $\boldsymbol{\mathcal{V}}$ to $\boldsymbol{\mathcal{Z}}$. To simplify notation, I retain the same symbols for these extended definitions. Restricting attention to the baseline domain with $\boldsymbol{\mathcal{V}}$ recovers the original objects defined in the main text.

A recommendation rule is *induced* by a scoring rule if it selects a seller with the highest non-negative score and no sellers if all scores are negative. A scoring rule is *optimal* if its induced recommendation rule is optimal.

This framework nests the baseline environment as a special case with $\boldsymbol{z} = \boldsymbol{v}$.

I now characterize an optimal recommendation mechanism in environments with an informed intermediary and with an uninformed intermediary.

Informed Intermediary. Suppose the intermediary observes match values, so $\boldsymbol{z} = \boldsymbol{v}$. Seller information does not affect the set of feasible recommendation rules, since sellers' private information is always θ_i which is equivalent to w_i conditioning on v_i. Thus the intermediary's problem coincides with the baseline mechanism design problem, with the only distinction arising from consumer behavior.

Under *allocative* choice, obedience is irrelevant; the optimal scoring rule is

$$\hat{s}_i^A(\boldsymbol{v}, \boldsymbol{w}) = w_i - \frac{1 - H(w_i \mid v_i)}{h(w_i \mid v_i)}.$$

Under *persuasive* choice, obedience introduces an adjustment term $\ell_i^*(\boldsymbol{v})$ as in (8). The optimal scoring rule is

$$\hat{s}_i^*(\boldsymbol{v}, \boldsymbol{w}) = w_i - \frac{1 - H(w_i \mid v_i)}{h(w_i \mid v_i)} + \ell_i^*(\boldsymbol{v}).$$

Uninformed Intermediary. Suppose the intermediary is uninformed, so $\boldsymbol{z} = \boldsymbol{\emptyset}$. By Relevance, sellers must then be fully informed. Since seller payoff depends on private information only through realized willingness to pay $w_i = w(v_i, \theta_i)$, and neither v_i nor θ_i can be elicited separately, the mechanism design problem reduces to a one-dimensional problem on $\boldsymbol{\mathcal{W}}$, with $(\hat{\boldsymbol{r}}, \hat{\boldsymbol{t}}) : \boldsymbol{\mathcal{W}} \to \Delta(\{0\} \cup \mathcal{N}) \times \mathbb{R}^N$.

By the envelope theorem, revenue equals expected unconditional virtual values,

$$\int_{\mathcal{W}} \sum_{i \in \mathcal{N}} \hat{\varphi}^U(w_i)\hat{r}_i(\boldsymbol{w})\boldsymbol{H}(d\boldsymbol{w}), \qquad \hat{\varphi}^U(w_i) = w_i - \frac{1 - H(w_i)}{h(w_i)}. \tag{12}$$

Define the allocation rule

$$\hat{\rho}_i^U(\boldsymbol{w}) = 1 \text{ if } \hat{\varphi}^U(w_i) > \max\left\{\max_{j \in \mathcal{N}\setminus\{i\}} \hat{\varphi}^U(w_j), 0\right\}$$

with $\overline{v}^U = \hat{v}_{i|i}(\hat{\boldsymbol{\rho}}^U)$ and $\underline{v}^U = \hat{v}_{i|0}(\hat{\boldsymbol{\rho}}^U)$.

If consumer choice is allocative, $\hat{\boldsymbol{\rho}}^U$ is optimal and the optimal scoring rule is

$$\hat{s}_i^U(\boldsymbol{w}) = w_i - \frac{1 - H(w_i)}{h(w_i)}. \tag{13}$$

If consumer choice is persuasive, the recommendation rule must satisfy obedience,

$$\int_{\mathcal{W}} (\hat{v}_i(w_i) - \hat{v}_j(w_j))r_i(\boldsymbol{w})\boldsymbol{H}(d\boldsymbol{w}) \geq 0 \text{ for all } i, j \in \{0\} \cup \mathcal{N},$$

where $\hat{v}_k(w_k) = \mathbb{E}(v_k \mid w_k)$ for $k \in \mathcal{N}$ and $\hat{v}_0 = v_0$. The optimal scoring rule is

$$\hat{s}_i^U(\boldsymbol{w}) = w_i - \frac{1 - H(w_i)}{h(w_i)} + \hat{\ell}_i^U(\boldsymbol{w}).$$

where

$$\hat{\ell}_i^U(\boldsymbol{w}) = \begin{cases} 0 & \text{if } v_0 \in [\underline{v}^U, \overline{v}^U] \\ \lambda_0^U(\hat{v}(w_i) - v_0) & \text{if } v_0 > \overline{v}^U \\ \lambda_1^U\left(\frac{1}{N}\sum_{k \in \mathcal{N}} \hat{v}(w_k) - v_0\right) & \text{if } v_0 < \underline{v}^U \end{cases}$$

where $\lambda_0^U, \lambda_1^U \geq 0$ are Lagrange multipliers on $OB_{0|i}$ and $OB_{i|0}$, respectively.

Strict Monotonicity and Single-Crossing.

Lemma 5. *Fix an environment. An optimal scoring rule $\hat{\boldsymbol{s}}$ satisfies:*

1. *(Strict Monotonicity) For every $i \in \mathcal{N}$, $\hat{s}_i$ is strictly increasing in w_i for all $(\boldsymbol{z}, \boldsymbol{w}_{-i}) \in \boldsymbol{\mathcal{Z}} \times \boldsymbol{\mathcal{W}}_{-i}$,*
2. *(Strict Single-Crossing Property) For every $i \neq j$ in $\mathcal{N}$, $\hat{s}_i - \hat{s}_j$ satisfies strict*

single-crossing in w_i: for any $x < x'$ in $\mathcal{W}$ and any $(\boldsymbol{z}, \boldsymbol{w}_{-i}) \in \boldsymbol{\mathcal{Z}} \times \boldsymbol{\mathcal{W}}_{-i}$,

$$\hat{s}_i(\boldsymbol{z}, x, \boldsymbol{w}_{-i}) \geq \hat{s}_i(\boldsymbol{z}, x, \boldsymbol{w}_{-i}) \implies \hat{s}_i(\boldsymbol{z}, x', \boldsymbol{w}_{-i}) > \hat{s}_i(\boldsymbol{z}, x', \boldsymbol{w}_{-i})$$

Proof. For an informed intermediary and a persuasive consumer,

$$\hat{s}_i(\boldsymbol{v}, \boldsymbol{w}) = w_i - \frac{1 - H(w_i \mid v_i)}{h(w_i \mid v_i)} + \ell_i^*(\boldsymbol{v}).$$

The virtual term is strictly increasing in w_i while $\ell_i^*(\boldsymbol{v})$ does not depend on w_i, so $\hat{s}_i$ is strictly increasing in w_i. For any $j \neq i$, $\hat{s}_j$ is independent of w_i, hence $\hat{s}_i - \hat{s}_j$ is strictly increasing in w_i, implying strict single-crossing property. The same argument applies for an allocative consumer.

For an uninformed intermediary and a persuasive consumer,

$$\hat{s}_i^U(\boldsymbol{w}) = w_i - \frac{1 - H(w_i)}{h(w_i)} + \hat{\ell}_i^U(\boldsymbol{w}).$$

The virtual term is strictly increasing in w_i while $\ell_i^U(\boldsymbol{w})$ is non-decreasing in w_i, so that $\hat{s}_i$ is strictly increasing in w_i. For any $j \neq i$, $\hat{\ell}_i^U(\boldsymbol{w}) - \hat{\ell}_j^U(\boldsymbol{w})$ is non-decreasing in w_i, so that $\hat{s}_i^U - \hat{s}_j^U$ is strictly increasing in w_i, implying strict single-crossing property. The same agument applies for an allocative consumer. □

Auction Rules. Let $\hat{\boldsymbol{w}} : \boldsymbol{\mathcal{Z}} \times \mathbb{R} \to \boldsymbol{\mathcal{W}}$ be an *inferred willingness to pay* rule, where each $\hat{w}_i(\boldsymbol{z}, b_i)$ denotes the intermediary's inferred willingness to pay for seller i given information $\boldsymbol{z}$ and bids b_i.[6]

Definition 2. Under a *weighted second inferred willingness to pay auction (wSIWA)* with $\hat{\boldsymbol{w}}$ and $\hat{\boldsymbol{s}}$, sellers submit bids b_i. Given $(\boldsymbol{z}, \boldsymbol{b})$, the intermediary computes inferred willingness to pays $\hat{\boldsymbol{w}}(\boldsymbol{z}, \boldsymbol{b})$, assigns scores $\hat{s}_i(\boldsymbol{z}, \hat{\boldsymbol{w}}(\boldsymbol{z}, \boldsymbol{b}))$, and selects the seller with the highest non-negative score (ties broken by a fixed rule). The winner pays the lowest willingness to pay at which she would still be selected.

Fix i, $\boldsymbol{z}$ and $\boldsymbol{b}_{-i}$. Let $\hat{\boldsymbol{w}}_{-i} = \hat{\boldsymbol{w}}_{-i}(\boldsymbol{z}, \boldsymbol{b}_{-i})$. Define

$$\hat{p}_i(\boldsymbol{z}, \hat{\boldsymbol{w}}_{-i}) = \inf\left\{x \in \mathcal{W} \mid \hat{s}_i(\boldsymbol{z}, x, \hat{\boldsymbol{w}}_{-i}) \geq \max\left\{0, \max_{j \neq i} \hat{s}_j(\boldsymbol{z}, x, \hat{\boldsymbol{w}}_{-i})\right\}\right\}, \tag{14}$$

[6]I restrict $\hat{w}_i$ to depend only on $(\boldsymbol{z}, b_i)$, reflecting that $\boldsymbol{b}_{-i}$ is not informative about (v_i, θ_i).

with $\hat{p}_i(\boldsymbol{z}, \hat{\boldsymbol{w}}_{-i}) = \sup \mathcal{W}$ if the set is empty. Seller i wins if and only if

$$\hat{w}_i(z_i, b_i) > \hat{p}_i(\boldsymbol{z}, \hat{\boldsymbol{w}}_{-i}),$$

and, conditional on winning, pays $\hat{p}_i(\boldsymbol{z}, \hat{\boldsymbol{w}}_{-i})$.

An auction rule *implements a recommendation rule* if there is a symmetric Bayes–Nash equilibrium such that the induced recommendation coincides almost surely with the recommendation rule. An auction rule *implements a scoring rule* if it implements a recommendation rule induced by the scoring rule. An auction rule *implements an optimal recommendation rule* if, for every admissible specification of primitives (w, F, G), it implements a recommendation rule for those primitives.

Let $\mathcal{Y} = \Theta \cup (\mathcal{V} \times \Theta)$ be a set of signals that each seller observes about states. Define each seller's information function $y : \mathcal{V} \times \Theta \to \mathcal{Y}$ as

$$y(v, \theta) = \begin{cases} \theta & \text{if sellers are partially informed} \\ (v, \theta) & \text{if sellers are fully informed} \end{cases}.$$

Lemma 6. *Fix an environment, an inferred willingness to pay $\hat{\boldsymbol{w}}$ and a scoring rule $\hat{\boldsymbol{s}}$. Suppose that for each seller i, there exists a bid function $b_i : \mathcal{Y} \to \mathbb{R}$ such that for all states $(\boldsymbol{v}, \boldsymbol{\theta})$, the intermediary's information $\boldsymbol{z} = \boldsymbol{z}(\boldsymbol{v})$ and all opponents' bids,*

$$\hat{w}_i(\boldsymbol{z}, b_i(y(v_i, \theta_i))) = w_i(v_i, \theta_i). \tag{15}$$

Assume $\hat{\boldsymbol{s}}$ be a scoring rule satisfies strict monotonicity and strict single-crossing property. Then, in the weighted second inferred willingness to pay auction with $\hat{\boldsymbol{w}}$ and $\hat{\boldsymbol{s}}$, bidding $b_i(y(v_i, \theta_i))$ is a weakly dominant strategy for seller i, and the auction implements $\hat{\boldsymbol{s}}$.

Proof. Fix i, $(\boldsymbol{v}, \boldsymbol{\theta})$, $(\boldsymbol{z}, \boldsymbol{b}_{-i})$, and hence, $\hat{\boldsymbol{w}}_{-i} = \hat{\boldsymbol{w}}_{-i}(\boldsymbol{z}, \boldsymbol{b}_{-i})$. By strict monotonicity and single-crossing property, the set of inferred values x such that

$$\hat{s}_i(\boldsymbol{z}, x, \hat{\boldsymbol{w}}_{-i}) \geq \max\left\{0, \max_{j \neq i} \hat{s}_j(\boldsymbol{z}, x, \hat{\boldsymbol{w}}_{-i})\right\}$$

is an upper interval in x. Hence there exists a cutoff $\hat{p}_i(\boldsymbol{z}, \hat{\boldsymbol{w}}_{-i})$ such that seller i wins if and only if $\hat{w}_i > \hat{p}_i(\boldsymbol{z}, \hat{\boldsymbol{w}}_{-i})$.

Let $w_i = w(v_i, \theta_i)$ and $p_i = \hat{p}(\boldsymbol{z}, \hat{\boldsymbol{w}}_{-i})$. The payoff for winning is $w_i - \hat{p}_i$; losing is 0. Thus, seller i prefers to win if and only if $w_i > \hat{p}_i$. By assumption, bidding $b_i = b_i(y(v_i, \theta_i))$ induces $\hat{w}_i(\boldsymbol{z}, b_i) = w_i$, implements exactly this choice, and therefore, is weakly dominant. Hence, the auction implements $\hat{\boldsymbol{s}}$. □

All auction formats considered below are special cases of wSIWA that differ only in how inferred willingness to pay $\hat{\boldsymbol{w}}$ and the scoring rule $\hat{\boldsymbol{s}}$ are constructed from bids and intermediary information.

Definition 3. The *weighted second maximum willingness to pay auction (wSMWA)* is the wSIWA with inferred willingness to pay rule

$$\hat{w}_i(\boldsymbol{z}, b_i) = \mathbb{E}\left(w(v_i, \theta_i) \mid \boldsymbol{z}, b_i = \sup_{y(v_i,\theta_i)=y_i} w(v_i, \theta_i) \right),$$

that is, the posterior expectation of seller i's realized valuation given its information $\boldsymbol{z}$, interpreting the bid as reporting the seller's maximum willingness to pay,

$$b_i(y_i) = \sup_{y(v_i,\theta_i)=y_i} w(v_i, \theta_i). \tag{16}$$

Unless stated otherwise, wSMWA uses the optimal scoring rule in each environment.

Proposition 4. *wSMWA is optimal for all environments.*

Proof. Fix an environment and let $\hat{\boldsymbol{s}}$ be the optimal scoring rule. By Lemma 5, $\hat{\boldsymbol{s}}$ satisfies strict monotonicity and strict single-crossing property. By Lemma 6, it suffices to show that condition (15) under the bid function $b_i(.)$ defined in (16).

Suppose sellers are partially informed, $y_i = \theta_i$. By Relevance, the intermediary is informed, so $\boldsymbol{z} = \boldsymbol{v}$. Under (16), $b_i(y_i) = \bar{w}(\theta_i)$, which is strictly increasing and therefore identifies θ_i. Together with $\boldsymbol{z} = \boldsymbol{v}$, this determines $w(v_i, \theta_i)$, so the posterior is degenerate and condition (15) holds.

Suppose sellers are fully informed, $y_i = (v_i, \theta_i)$. Under (16), $b_i(y_i) = \sup_{y(v_i,\theta_i)=y_i} w(v_i, \theta_i) = w(v_i, \theta_i)$, so again condition (15) holds.

In either case, by Lemma 6, wSMWA implements $\hat{\boldsymbol{s}}$. □

Definition 4. The *weighted second maximum willingness to pay auction (wSMWA)* is the wSIWA with inferred willingness to pay rule

$$\hat{w}_i(\boldsymbol{z}, b_i) = b_i,$$

that is, interpreting each bid as the seller's true willingness to pay. Unless stated otherwise, wSPA uses the optimal scoring rule in each environment.

An auction format *reduces to* another format if their recommendation and payment rules coincide for every realization of the intermediary's information $\boldsymbol{z}$ and bids $\boldsymbol{b}$; when referring to an auction format, I implicitly mean the optimal member of the corresponding class.

Proposition 5. *wSMWA reduces to wSPA if and only if sellers are fully informed.*

Proof. The two auctions use the same scoring and payment rule, it suffices to compare

$$\hat{w}_i^{wSPA}(\boldsymbol{z}, b_i) = b_i, \quad \hat{w}_i^{wSMWA}(\boldsymbol{z}, b_i) = \mathbb{E}\left(w(v_i, \theta_i) \mid \boldsymbol{z}, b_i = \sup_{y(v_i,\theta_i)=y_i} w(v_i, \theta_i) \right).$$

If sellers are fully informed, the supremum equals $w(v_i, \theta_i)$, so $\hat{w}_i^{wSMWA}(\boldsymbol{z}, b_i) = b_i = \hat{w}_i^{wSPA}(\boldsymbol{z}, b_i)$ for all $(\boldsymbol{z}, \boldsymbol{b})$. Scores, hence outcomes, coincide. Reduction holds.

If partially informed, $\boldsymbol{z} = \boldsymbol{v}$ and the supremum equals $\bar{w}(\theta_i)$, so $\hat{w}_i^{wSMWA}(\boldsymbol{z}, b_i) = w(v_i, \bar{w}^{-1}(b_i))$, which varies with v_i while $\hat{w}_i^{wSPA}(\boldsymbol{z}, b_i) = b_i$ do not. Hence, there exists $(\boldsymbol{z}, \boldsymbol{b})$ for which score rankings differ, so reduction fails. □

Proposition 6. *wSPA is optimal if and only if sellers are fully informed.*

Proof. If sellers are fully informed, wSPA coincides with wSMWA by Proposition 5, and therefore, is optimal.

Suppose sellers are partially informed. By Relevance, the intermediary is informed. Consider two sellers $i \in \{1, 2\}$ with $w(v_i, \theta_i) = v_i + \theta_i$ with v_i, θ_i i.i.d. $U[0, 1]$ which satisfies Assumption 1. Let $v_0 = \mathbb{E}(v_i)$, so the same scoring rule applies to both persuasive and allocative choices. The optimal scoring rules are

$$s_i^*(\boldsymbol{v}, \boldsymbol{\theta}) = 2\theta_i - 1 + v_i, \qquad \hat{s}_i^*(\boldsymbol{v}, \boldsymbol{w}) = 2w_i - 1 - v_i$$

so s_i^* is strictly increasing in v_i while $\hat{s}_i^*$ is strictly decreasing in v_i.

Suppose wSPA is optimal. Then, it admits a strictly increasing symmetric equilibrium bid $b(\theta)$, so b is continuous except on a countable set; hence, for any $\epsilon > 0$, there is $\theta^* \in (1 - \epsilon, 1)$ at which b is continuous. Consider

$$(v_1, \theta_1) = (1, \theta^*), \qquad (v_2, \theta_2) = (0, \theta^*).$$

Under the optimal rule, bidder 1 is selected:

$$s_1^*(\boldsymbol{v}, \boldsymbol{\theta}) = 2\theta^* > 2\theta^* - 1 = \max(s_2^*(\boldsymbol{v}, \boldsymbol{\theta}), 0)$$

Under wSPA, bidder 1 is not selected:

$$\hat{s}_1^{wSPA}(\boldsymbol{v}, \boldsymbol{b}(\boldsymbol{\theta})) = 2b(\theta^*) - 2 < 2b(\theta^*) - 1 = \hat{s}_2^{wSPA}(\boldsymbol{v}, \boldsymbol{b}(\boldsymbol{\theta})).$$

Because the inequalities are strict and the scoring rules are continuous at $(1, 0, \theta^*, \theta^*)$, they hold on an open neighborhood of this point and thus occur with positive probability. Hence, wSPA fails to implement the optimal rule and is not optimal.

□

Definition 5. The second-price auction (SPA) is the wSIWA with

$$\hat{w}_i(\boldsymbol{z}, b_i) = b_i,$$

and with scoring rule given by a strictly increasing function of willingness to pay,

$$\hat{\boldsymbol{s}}_i(\boldsymbol{z}, \boldsymbol{w}) = s^\dagger(w_i).$$

Equivalently, SPA selects the seller with the highest bid and charges the second highest bid with reserve price $\kappa = (s^\dagger)^{-1}(0)$.

Proposition 7. *wSPA reduces to SPA if the environment is auction-like.*

Proof. By definition, SPA is a special case of wSPA in which bidders are ranked by a strictly increasing function of their bids. Therefore, wSPA reduces to SPA if and only if its scoring rule is a strictly increasing function of each bidder's bid alone.

If the environment is auction-like, the optimal scoring rule is (B.3) with $w_i = b_i$. This is strictly increasing in each bidder's bid alone, so wSPA reduces to SPA.

If the environment is not auction-like, then either (i) the intermediary is informed or (ii) consumer choice is persuasive. In either case, the wSPA score is not a function of each bidder's bid alone, and therefore, reduction fails. □

Proposition 8. *SPA is optimal if and only if the environment is auction-like.*

Proof. Auction-like environments collapse to Myerson (1981) where SPA is optimal.

Suppose the environment is not auction-like.

Case 1. Suppose the consumer's choice is persuasive, but other assumptions remain the same. Consider two sellers $i \in \{1, 2\}$ with $w(v_i, \theta_i) = v_i + \theta_i$ with v_i, θ_i i.i.d. $U[0, 1]$ which satisfies Assumption 1, and $v_0 < \underline{v}^*$. The optimal scoring rule

$$\hat{s}_i^U(\boldsymbol{w}) = \varphi^U(w_i) + \lambda_1 \left(\frac{1}{N} \sum_{k \in \mathcal{N}} \frac{1}{2} w_k - v_0 \right).$$

is continuous and satisfies $\inf \varphi^U(\boldsymbol{w}) < 0 < \sup \varphi^U(\boldsymbol{w})$. Let $\alpha^U(w) = \varphi^U(w) + \lambda (w - v_0)$, crossing zero at a unique $w^* \in (0, 2)$. Fix small $\epsilon > 0$ and let

$$\boldsymbol{w} = (w^* + \epsilon, w^*, ..., w^*), \qquad \boldsymbol{w}' = (w^* + \epsilon, 0, ..., 0).$$

SPA depends only on the ranking of bids (equivalently, the ranking of w_i) and therefore treats $\boldsymbol{w}$ and $\boldsymbol{w}'$ identically: seller 1 is selected in both states or in neither. In contrast, under the optimal score all sellers' scores are positive at $\boldsymbol{w}$, so seller 1 is selected, while at $\boldsymbol{w}'$ the aggregate term is negative that $\hat{s}_i^U(\boldsymbol{w}') < 0$ for sufficiently small ϵ, so seller 1 is not selected. By strict inequality and continuity, this disagreement occurs on a set of positive probability; SPA cannot implement the optimal rule.

Case 2. Suppose the intermediary is informed and sellers are fully informed. Use the same primitives as Case 1 but set $v_0 = \mathbb{E}(v_i)$, so the same scoring rule applies to both persuasive and allocative choices. The optimal scoring rule $\mathbf{s}^*$ in $(\boldsymbol{v}, \boldsymbol{\theta})$-space is

$$s_i^*(\boldsymbol{v}, \boldsymbol{\theta}) = 2\theta_i - 1 + v_i.$$

Consider $(v_i, \theta_i) = (0, 1)$ and $(v_j, \theta_j) = (1, 1 - c)$ for some $c \in (1/2, 1)$. Then

$$s_i^*(\boldsymbol{v}, \boldsymbol{\theta}) = 1 > 1 - 2c = \max(s_j^*(\boldsymbol{v}, \boldsymbol{\theta}), 0),$$

so the optimal rule selects i. But

$$w(v_i, \theta_i) = 1 < 2 - c = w(v_j, \theta_j),$$

so SPA does not select i. Again, by continuity this disagreement holds on a set of positive probability. Hence SPA cannot implement the optimal rule.

Case 3. Suppose the intermediary is informed and sellers are partially informed. The optimal scoring rule depends on $\boldsymbol{v}$. In SPA, bids, and thus scores, depend only on $\boldsymbol{\theta}$; SPA cannot implement the optimal rule. This completes the proof. $\square$

C Proofs for Section 5 and Section 6

C.1 Proof for Theorem 4

Let $d \in \{0\} \cup \mathcal{N}$ be the default and $i \in \mathcal{N}$ be a product. Since $\boldsymbol{r}^*$ is symmetric,

$$CW(\boldsymbol{r}^*) - CW^D = \left(\hat{v}_{i|i}(\boldsymbol{r}^*) - \hat{v}_{d|i}(\boldsymbol{r}^*)\right)\left(1 - \Pr\left(r_0^*\right)\right) + \left(v_0 - \hat{v}_{d|i}(\boldsymbol{r}^*)\right)\Pr\left(r_0^*\right).$$

If $v_0 \geq \overline{v}^*$, then $d = 0$, $OB_{0|i}$ binds $\hat{v}_{i|i}(\boldsymbol{r}^*) = v_0$, and $CW(\boldsymbol{r}^*) = v_0 = CW^D$.

If $\underline{v}^* < v_0 < \overline{v}^*$, then: if $d = 0$, then $\hat{v}_{i|i}(\boldsymbol{r}^*) = \overline{v}^* > v_0 = \hat{v}_{d|i}(\boldsymbol{r}^*)$; if $d \in \mathcal{N}$, then $\hat{v}_{0|0}(\boldsymbol{r}^*) = v_0 > \underline{v}^* = \hat{v}_{d|0}(\boldsymbol{r}^*)$. In either case, $CW(\boldsymbol{r}^*) > CW^D$.

If $v_0 \leq \underline{v}^*$, then $d \in \mathcal{N}$, $OB_{d|i}$ is slack for any $i \in \mathcal{N} \setminus \{d\}$, $\hat{v}_{i|i}(\boldsymbol{r}^*) > \hat{v}_{d|i}(\boldsymbol{r}^*)$, and $CW(\boldsymbol{r}^*) > CW^D$.

C.2 Proof for Theorem 5

Assumption 2 is maintained throughout this subsection. Since $w \geq 0$, the outside option is recommended with probability 0 under the optimal allocation rule

$$\rho_i^S(\boldsymbol{v}, \boldsymbol{\theta}) = 1 \text{ if } w(v_i, \theta_i) > \max_{j \in \mathcal{N} \setminus \{i\}} w(v_j, \theta_j).$$

Let $\overline{v}^S = \hat{v}_{i|i}(\boldsymbol{\rho}^S) \geq \mathbb{E}(v_i)$. If $v_0 \leq \overline{v}^S$, obedience is slack, so $\boldsymbol{r}^S = \boldsymbol{\rho}^S$ and $CW(\boldsymbol{r}^S) = \overline{v}^S$; if $v_0 > \overline{v}^S$, $OB_{0|i}$ binds and $CW(\boldsymbol{r}^S) = v_0$. Hence, $CW(\boldsymbol{r}^S) = \max(\overline{v}^S, v_0)$.

Case 1. $v_0 \geq \max(\overline{v}^*, \overline{v}^S)$**:** $OB_{0|i}$ binds in both, hence $CW(\boldsymbol{r}^*) = v_0 = CW(\boldsymbol{r}^S)$.

Case 2. $v_0 \leq \underline{v}^*$**:** Introduce the auxiliary allocation rule

$$\rho_i^O(\boldsymbol{v}, \boldsymbol{\theta}) = 1 \text{ if } \varphi(v_i, \theta_i) > \max_{j \in \mathcal{N} \setminus \{i\}} \varphi(v_j, \theta_j),$$

maximizing virtual willingness to pay. Let $\overline{v}^O = \hat{v}_{i|i}(\boldsymbol{\rho}^O)$.

Lemma 7. *Let $v_0 \leq \underline{v}^*$. Then, $CW(\boldsymbol{r}^*) < \overline{v}^O$.*

Proof. Since $\boldsymbol{\rho}^O$ always recommends a product, $CW(\boldsymbol{\rho}^O) = \overline{v}^O$. Since $\boldsymbol{r}^* \neq \boldsymbol{\rho}^O$ only when $r_0^* = 1$,

$$\overline{v}^O - CW(\boldsymbol{r}^*) = \int_{\boldsymbol{\mathcal{V}} \times \boldsymbol{\Theta}} \sum_{i \in \mathcal{N}} (v_i - v_0) \rho_i^O(\boldsymbol{v}, \boldsymbol{\theta}) r_0^*(\boldsymbol{v}, \boldsymbol{\theta}) \boldsymbol{F}(d\boldsymbol{v}) \boldsymbol{G}(d\boldsymbol{\theta}).$$

Applying $\hat{v}_{i|0}(\boldsymbol{r}^*) = v_0$ because $OB_{i|0}$ binds under $\boldsymbol{r}^*$ and symmetry, this reduces to

$$\int_{v > v'} (v - v') \big(p_i(v, v') - p_i(v', v) \big) F(dv) F(dv') > 0,$$

where

$$p_i(v_i, v_j) = \int_{\boldsymbol{\mathcal{V}}_{-ij} \times \boldsymbol{\Theta}} \rho_i^O(\boldsymbol{v}, \boldsymbol{\theta}) r_0^*(\boldsymbol{v}, \boldsymbol{\theta}) \boldsymbol{F}(d\boldsymbol{v}_{-ij}) \boldsymbol{G}(d\boldsymbol{\theta}).$$

It therefore suffices to show

$$p_i(v, v') > p_i(v', v) \quad \text{for all } v > v'.$$

Fix $(\boldsymbol{v}_{-ij}, \boldsymbol{\theta}_{-ij})$ and $v_i = v > v' = v_j$. Let

$$X = \varphi(v, \theta), \quad Y = \varphi(v', \theta'), \qquad a = \max_{k \in \mathcal{N} \setminus \{i,j\}} \varphi(v_k, \theta_k), \qquad b = \lambda_1 \left(\frac{1}{N} \sum_{k \in \mathcal{N}} v_k - v_0 \right).$$

Then,

$$\rho_i^O = 1 \iff X > \max(Y, a), \qquad r_0^* = 1 \iff \max(X, Y, a) < b.$$

Hence, $\rho_i^O r_0^* = 1 \iff \{X > \max(Y, a)\} \cap \{\max(X, Y, a) < b\}$. Let $C = \{a < \max(X, Y) < b\}$. Since $X > \max(Y, a)$ implies $\max(X, Y, a) = X$, this reduces to

$$\rho_i^O r_0^* = 1 \iff \{X > Y\} \cap C.$$

Similarly, when $v_i = v' < v = v_j$, $\rho_i^O r_0^* = 1 \iff \{X < Y\} \cap C$. Therefore,

$$p_i(v, v' \mid \boldsymbol{v}_{-ij}, \boldsymbol{\theta}_{-ij}) = \Pr(X > Y, C), \qquad p_i(v', v \mid \boldsymbol{v}_{-ij}, \boldsymbol{\theta}_{-ij}) = \Pr(X < Y, C).$$

Let $\tilde{H}_Z$ and $\tilde{h}_Z$ be the cdf and density of $Z \in \{X, Y\}$. Then,

$$\Pr(X > Y, C) = \Pr(X > Y, a < X < b) = \int_a^b \Pr(X > Y \mid X = t)\tilde{h}_X(t)dt = \int_a^b \tilde{H}_Y(t)\tilde{h}_X(t)dt$$

Similarly, $p_i(v', v \mid \boldsymbol{v}_{-ij}, \boldsymbol{\theta}_{-ij}) = \int_a^b \tilde{H}_X(t)\tilde{h}_Y(t)dt$. Therefore,

$$\Delta = p_i(v, v' \mid \boldsymbol{v}_{-ij}, \boldsymbol{\theta}_{-ij}) - p_i(v', v \mid \boldsymbol{v}_{-ij}, \boldsymbol{\theta}_{-ij}) = \int_a^b \tilde{H}_X(t)\tilde{H}_Y(t)\left(\frac{\tilde{h}_X(t)}{\tilde{H}_X(t)} - \frac{\tilde{h}_Y(t)}{\tilde{H}_Y(t)}\right)dt.$$

Strict log-concavity of G and $m_\theta^\varphi(v, \theta) \leq 0$ imply $\frac{\tilde{h}_X(t)}{\tilde{H}_X(t)} > \frac{\tilde{h}_Y(t)}{\tilde{H}_Y(t)}$ for all t, so $\Delta > 0$. Integrating over $(\boldsymbol{v}_{-ij}, \boldsymbol{\theta}_{-ij})$ gives $p_i(v, v') > p_i(v', v)$, completing the proof.

□

Lemma 8. $\overline{v}^O < \overline{v}^S$.

Proof. Since $m^w > m^\varphi$, the ranking under w is more responsive to match values than that under φ. Thus the maximizer under w has weakly higher match value, strictly with positive probability, implying $\overline{v}^O < \overline{v}^S$. □

Therefore, for $v_0 \leq \underline{v}^*$, by Lemma 7 and Lemma 8,

$$CW(\boldsymbol{r}^*) < \overline{v}^O \leq \overline{v}^S = CW(\boldsymbol{r}^S).$$

Case 3. $v_0 \in (\underline{v}^*, \max(\overline{v}^*, \overline{v}^S))$**:** If $\overline{v}^S \geq \overline{v}^*$, then,

$$CW(\boldsymbol{r}^*) = \overline{v}^*\big(1 - \Pr(r_0^*)\big) + v_0 \Pr(r_0^*) < \overline{v}^* \leq \overline{v}^S = CW(\boldsymbol{r}^S).$$

If $\overline{v}^S < \overline{v}^*$, $CW(\boldsymbol{r}^*)$ is affine in v_0, equals $CW(\boldsymbol{r}^S)$ at $v_0 = \overline{v}^*$, and is strictly lower for $v_0 \leq \underline{v}^*$; hence, it crosses $\underline{v}^S \in (\underline{v}^*, \overline{v}^*)$ once from below at a unique $\overline{v}^S \in (\underline{v}^*, \overline{v}^*)$.

C.3 Proof for Example 3

Assumptions in Section 2: Let $x = av + b\theta$ and $w(v, \theta) = c(x)$. Since $a, b, c' > 0$, $w_v, w_\theta > 0$. From $\varphi = c + c'b\psi$, $\varphi_v = a(c' - c''b\psi) > 0$ since $c'' \leq 0$ and $\psi \geq 0$, and $\varphi_\theta(v, \theta) = c'b - c''(b)^2\psi - c'b\psi' > 0$ because $\psi' < 0$.
Assumption 1: Under log-concavity, $x = av + b\theta$ has a log-concave density, hence decreasing inverse hazard rate. Since $w = c(x)$ is strictly increasing, the same holds

for w, implying φ^U is strictly increasing. Log-concavity of x further implies monotone likelihood ratio in v, which is preserved by the monotone transformation $w = c(x)$.
Assumption 2: Log-concavity of g implies that of G. The inequality $m^w \geq m^\varphi$ is equivalent to $\psi' \leq 0$, which holds by assumption. Moreover, $(m^\varphi)_\theta < 0$ is equivalent to $(1 + k\psi)\psi'' - \psi'(k_\theta\psi + k\psi) < 0$, where $k = -b\frac{c''}{c'}$. This follows from $\psi \geq 0, \psi' < 0, \psi'' < 0, k \geq 0$ and $k' > 0$, where the last holds because $(\log c')'' > 0$.

References

Admati, A. R. and P. Pfleiderer (1990): "Direct and Indirect Sale of Information," *Econometrica*, 58, 901–928.

Athey, S. and G. Ellison (2011): "Position Auctions with Consumer Search," *The Quarterly Journal of Economics*, 126, 1213–1270.

Bergemann, D. and A. Bonatti (2024): "Data, Competition, and Digital Platforms," *American Economic Review*, 114, 151–191.

Bergemann, D., A. Bonatti, and A. Smolin (2018): "The Design and Price of Information," *American Economic Review*, 108, 1–48.

Bergemann, D., A. Bonatti, and N. Wu (2025): "How Do Digital Advertising Auctions Impact Product Prices?" *The Review of Economic Studies*, 92, 2330–2358.

Bergemann, D. and S. Morris (2017): "Information Design: A Unified Perspective," Cowles Foundation Discussion Paper 2075R, Cowles Foundation for Research in Economics, Yale University.

——— (2019): "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57, 44–95.

Crémer, J. and P. R. McLean (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257.

Dworczak, P. (2020): "Mechanism Design With Aftermarkets: Cutoff Mechanisms," *Econometrica*, 88, 2629–2661.

Dworczak, P. and A. Kolotilin (2019): "The Persuasion Duality," *SSRN Electronic Journal*.

Dworczak, P. and G. Martini (2019): "The Simple Economics of Optimal Persuasion," *Journal of Political Economy*, 56.

Edelman, B., M. Ostrovsky, and M. Schwarz (2007): "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97, 18.

Evans, R. and I.-U. Park (2026): "Selling Information for Bilateral Trade," Working Paper, R&R at Theoretical Economics.

Galperti, S. and J. Perego (2018): "A Dual Perspective on Information Design," *Working Paper*.

Gentzkow, M. and E. Kamenica (2016): "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *American Economic Review*, 106, 597–601.

Gomes, R. (2014): "Optimal Auction Design in Two-Sided Markets," *The RAND Journal of Economics*, 45, 248–272.

Ichihashi, S. and A. Smolin (2025a): "Buyer-Optimal Algorithmic Recommendations," Tech. Rep. 2309.12122, arXiv, working paper, revised Jun 2025.

——— (2025b): "Data provision to an informed seller," *Games and Economic Behavior*, 153, 131–144.

Inderst, R. and M. Ottaviani (2012): "Competition through Commissions and Kickbacks," *American Economic Review*, 102, 780–809.

Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Kolotilin, A. (2018): "Optimal Information Disclosure: A Linear Programming Approach," *Theoretical Economics*, 13, 607–635.

Kolotilin, A., R. Corrao, and A. Wolitzky (2025): "Persuasion and Matching: Optimal Productive Transport," *Journal of Political Economy*, 133, 1334–1381.

McAfee, R. P. and J. P. Reny (1992): "Correlated Information and Mecanism Design," *Econometrica*, 60, 395–421.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

——— (1982): "Optimal Coordination Mechanisms in Generalized Principal–Agent Problems - ScienceDirect," *Journal of Mathematical Economics*, 10.

——— (1983): "Bayesian Equilibrium and Incentive-Compatibility: An Introduction," *Working Paper*.

Nocke, V. and P. Rey (2024): "Consumer Search, Steering, and Choice Overload," *Journal of Political Economy*, 132, 1684–1739.

Rayo, L. and I. Segal (2010): "Optimal Information Disclosure," *Journal of*

Political Economy, 118, 949–987.

Segura-Rodriguez, C. (2021): "Selling Data," *Working Paper*.

Smolin, A. and T. Yamashita (2025): "Information Design in Smooth Games," Tech. Rep. 2202.10883, Working Paper, working paper, revised version July 2025.

Teh, T. and J. Wright (2022): "Intermediation and Steering: Competition in Prices and Commissions," *American Economic Journal: Microeconomics*, 14, 281–321.

Varian, H. R. (2007): "Position Auctions," *International Journal of Industrial Organization*, 25, 1163–1178.

Yang, K. (2024): "Equivalence in Business Models for Informational Intermediaries," Tech. rep., Working Paper.

Yang, K. H. (2021): "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences," *Working Paper*.