MARKET-MINDED INFORMATIONAL INTERMEDIARY AND UNINTENDED WELFARE LOSS

By

Wenji Xu and Kai Hao Yang

January 2022

COWLES FOUNDATION DISCUSSION PAPER NO. 2321

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Market-Minded Informational Intermediary and Unintended Welfare Loss*

Wenji Xu† Kai Hao Yang‡

January 25, 2022

Abstract

This paper examines the welfare effects of informational intermediation. A (short-lived) seller sets the price of a product that is sold through a (long-lived) informational intermediary. The intermediary can disclose information about the product to consumers, earns a fixed percentage of the sales revenue in each period, and has concerns about its prominence—the market size it faces in the future, which in turn is increasing in past consumer surplus. We characterize the Markov perfect equilibria and the set of subgame perfect equilibrium payoffs of this game and show that when the *market feedback* (i.e., how much past consumer surplus affects future market sizes) increases, welfare may decrease in the Pareto sense.

KEYWORDS: Informational intermediary, market size, market feedback, consumer surplus, Pareto-inferior outcomes, Markov perfect equilibrium, subgame perfect equilibrium.

JEL CLASSIFICATION: C73, D61, D82, D83, L15, M37

*This paper was previously circulated under the title: "Information Design in Bilateral Trade and Unintended Welfare Loss". This version supersedes the previous ones. We thank Pak Hung Au, James Bergin, Ben Brooks, Daniele Condorelli, Hector Chade, Yeon-Koo Che, Yi-Chun Chen, Joyee Deb, Jack Fanning, Nima Haghpanagh, Marina Halac, Wei He, Johannes Hörner, Emir Kamenica, Navin Kartik, Elliot Lipnowski, Yunan Li, Wooyoung Lim, Qingming Liu, Aniko Öry, Jacopo Perego, John Quah, Doron Ravid, Phil Reny, Hamid Sabourian, Evan Sadler, Tomasz Sadzik, Larry Samuelson, Anna Sanktjohasner, Philipp Strack, Nancy Stokey, Mark Whitmeyer, Yu Fu Wong, Qinggong Wu, Chen Zhao, and Jidong Zhou for their valuable comments. We also thank participants of various seminars and conferences at which this paper was presented. All errors are our own.

†College of Business, City University of Hong Kong, Email: wenjixu@cityu.edu.hk

‡School of Management, Yale University, Email: kaihao.yang@yale.edu

1 Introduction

In many markets, products are sold through an intermediary, who facilitates trade and provides product information to consumers before they make purchasing decisions. For example, in securities markets, financial advisors serve as intermediaries through which issuers sell securities to investors. At the same time, they provide information about financial products to investors. In insurance markets, insurance companies collaborate with insurance brokers who then persuade customers to buy these companies' insurance plans by providing more detailed information. Similarly, in the emerging online market, intermediaries such as influencers and key opinion leaders (KOL) provide product information to their followers, who then use that information to make purchase decisions. These *informational intermediaries* participate in the market by connecting product sellers with potential customers to whom they disclose product information. Moreover, such intermediaries often share two common features: (i) one of their common sources of income is through retaining a certain percentage of sales revenue as commission fees from product sellers;[1] and (ii) they have concerns about their own prominence, which evolves over time and depends on past consumers' satisfaction (e.g., positive experiences would lead to better reviews and more referral, which in turn make the intermediary more prominent in the future). Since an intermediary's prominence reflects the size of the market it faces, when interacting with a product seller, a forward-looking and revenue-maximizing informational intermediary would be *market-minded*. After all, the intermediary's long-term profit depends not only on sales revenues but also on its own prominence, which in turn is affected by consumers' satisfaction. Consequently, the level of *market feedback*—the degree to which consumers' satisfaction affect the intermediary's future prominence—would be crucial for understanding the intermediary's incentives, as well as welfare outcomes.[2] Higher level of market feedback means consumers' satisfaction

[1] While it is more common in practice that traditional intermediaries like insurance brokers and financial brokers are paid through commissions, such arrangements in the market for influencers may be seemingly less frequent. Nevertheless, one of the major income sources for online influencers are *affiliated links*. Namely, product sellers would partner with third party websites such as rStyle or ShopStyle, and then invite influencers to present information about their products to consumers while sharing the affiliated link generated by these third parties. These affiliated links can keep track of where each purchase is directed from, and the influencers will then be paid a fixed share of the revenue per purchase. For our purpose, this business model can also be thought of as a revenue-sharing arrangement between sellers and influencers.

[2] The level of market feedback can be affected by various factors. For example, market feedback level may be higher if the underlying communicative network among potential customers is more connected (i.e., the degree to which consumers engage in word-of-mouth after purchasing); if past customers have a better channel to leave reviews and deliver them to future customers more transparently (i.e., better rating system and better recommendation algorithm); if more consumers have alternatives to an intermediary (i.e., more intense competition among intermediaries); or if intermediaries are more specialized in information provision and provide fewer other services.

is more consequential to the intermediary's future prominence, and hence would lead to a more market-minded intermediary who assigns more weights to consumers' satisfaction when disclosing information. In contrast, lower level of market feedback means consumers' satisfaction has little impact on the intermediary's future prominence, and hence the intermediary is driven more by short-term revenue rather then consumers' satisfaction.

This paper studies how market feedback affects welfare in the presence of a market-minded informational intermediary. Specifically, since the level market feedback directly translates to the degree to which an informational intermediary is market-minded, a natural question follows: Does higher level market feedback always benefit consumers? Or more generally, how would market feedback affect welfare outcomes of the entire economy? In this paper, we show that not only is it possible that higher market feedback does not benefit the consumers, welfare might even *decrease* in the Pareto sense as the level of market feedback increases—leading to *unintended welfare losses*.

Specifically, we consider a dynamic game in discrete time with a (sequence of) short-lived seller (he) and a long-lived informational intermediary (she). In each period, a mass of short-lived consumers arrive. These consumers have unit demands and differentiated values, but do not observe their own values upon arrival. The seller first posts a price for his product. After observing the price, the intermediary then discloses information about the product to consumers so that they receive signals about their own values. With the information provided by the intermediary, each consumer then decides whether to buy the product. Lastly, the intermediary retains a fixed percentage of the sales revenue as commission fees. To model the feature that the intermediary's prominence (i.e., the mass of consumers who arrived in each period) depends on past customers' satisfaction, we assume that the mass of consumers arriving in each period (market size) depends on past consumer surplus. Specifically, we assume that market growth rate in each period is a nondecreasing affine function of the average consumer surplus in that period, with its slope being interpreted as the level of market feedback. As a result, the market size in period $t + 1$ onward would be larger if the average consumer surplus is higher in period t, and this effect is stronger when the market feedback is higher.

The result of unintended welfare loss is established by completely characterizing the equilibrium outcomes in our model. We first restrict attention to stationary-Markov perfect equilibrium outcomes (Theorem 1 and Theorem 2). When the market feedback level is low, there is a unique stationary-Markov perfect equilibrium outcome where the intermediary provides no information to consumers and the seller fully extracts all the surplus by charging a price that equals to the expected value. In this case, the consumer surplus is zero. As the market feedback level increases but remains below a certain threshold, the intermediary would begin to provide some information to consumers. The information provided is such that all consumers with values above a (nonzero) cutoff have the same interim expected value, which in turn equals to the price posted by the seller. As a result, consumer surplus remains zero and sales revenue becomes lower than the expected value. Moreover, as the market feedback increases, the equilibrium price increases and the sales revenue decreases, leading to Pareto inferior outcomes.

In addition to stationary-Markov equilibrium outcomes, we characterize the set of subgame perfect equilibrium payoffs for any fixed discount factor and market feedback level, with (Theorem 3) and without (Theorem 4) a restriction that the intermediary's present discounted profits are finite at every history. In both cases, when the market feedback level is below a certain threshold, there exists a subgame perfect equilibrium outcome that is Pareto-dominated by all other outcomes. Moreover, similar to the stationary-Markov perfect equilibrium outcomes, this Pareto-worst outcome becomes even worse (in the Pareto sense) as market feedback increases.

Together, the main characterizations suggest that higher market feedback may not benefit the consumers, and can even lead to Pareto inferior outcomes. This unintended welfare loss stems from the intrinsic structure of markets with an informational intermediary. After all, when an intermediary interacts with a product seller, the intermediary's incentives to bolster her future prominence creates difference in interest between her and the seller. On one hand, since the seller only interacts with the intermediary in the short term, the intermediary's future prominence is irrelevant to him and thus he only seeks to maximize current sales revenue. On the other hand, the intermediary is forward-looking and thus her profit depends on both current sales revenue and her future prominence, and the latter in turn depends on current consumer surplus. As a result, in each period and at any given price, the intermediary would prefer steering more consumers with values below the price to not purchasing compared to what the seller would prefer. Anticipating this, the seller would then charge a higher price to minimize the mass of consumers steered away by the intermediary, and hence decreasing efficiency.

Consequently, our results serve as a cautionary tale when it comes to improvements of market feedback. Increase in the level of market feedback (e.g., due to changes of market structures or technologies) may be detrimental for the entire market even though it renders a more market-minded intermediary. Meanwhile, policies that seek to improve market feedback (e.g., improving recommender systems or encouraging consumers' word-of-mouth behaviors)—or, from the same regard, policies that aim to incentivize intermediaries to enhance consumer surplus—should be evaluated and implemented carefully. After all, higher market feedback may be unambiguously undesirable when revenue sharing is the main business model between sellers and intermediaries and when these collaborating relationships are relatively short-term. To address these concerns, we show in Section 6 that such unintended welfare losses can be avoided either by transforming the revenue sharing business model to a subscription model or a customer-base monetization model, by imposing a cap on the market growth rate, or by establishing a long-term relationship between the product seller and the intermediary.

The rest of this paper is organized as follows. Section 2 below discusses the related literature. Section 3 introduces the model. We then characterize the stationary-Markov perfect equilibrium outcomes in Section 4, as well as subgame perfect equilibrium payoffs in Section 5. Section 6 and Section 7 discuss several policy implications and extensions, respectively. Section 8 then concludes.

2 Related Literature

This paper is related to the rapidly growing literature on information design and pricing, which introduces the design of information to various market structures, including consumer search (e.g., Anderson and Renault (2006), Board and Lu (2018), Au and Whitmeyer (2021), and Bergemann, Brooks, and Morris (2021)); monopoly (e.g., Bergemann, Brooks, and Morris (2015), Roesler and Szentes (2017), Du (2018), Ravid, Roseler, and Szentes (forthcoming), and Libgober and Mu (2021));[3] oligopoly (e.g., Boleslavsky, Hwang, and Kim (2019), Armstrong and Zhou (forthcoming); and Elliot, Galeotti, Koh, and Li (2021)); auctions (e.g., Bergemann and Pesendorfer (2007), Shi (2012), Bergemann, Brooks, and Morris (2017), Chen and Yang (2020), Kim and Koh (2021), Brooks and Du (2021), and Terstiege and Wasser (forthcoming)); and third-party intermediation (e.g., Yang (2019) and Yang (forthcoming)).

Among the aforementioned papers, our model is closer to Roesler and Szentes (2017), Ravid, Roseler, and Szentes (forthcoming), Du (2018), Libgober and Mu (2021), and Yang (2019). Similar to our model, all these papers involve designs of consumers' information structures in a monopolistic pricing setting. Specifically, Roesler and Szentes (2017) characterizes the surplus-maximizing information for consumers in a monopolistic setting, where the monopolist always posts a price optimally based on the information structure consumers have. Ravid, Roseler, and Szentes (forthcoming) characterizes the equilibrium outcomes when the monopolist posts a price and consumers acquire information simultaneously. Du (2018) solves for a robust selling mechanism (i.e., randomized posted price) that maximizes the seller's revenue guarantee, where the Nature always chooses the worst case consumer information given the selected mechanism. Libgober and Mu (2021) considers a dynamic pricing problem where consumers can delay purchases and the seller can commit to a price path, and the Nature chooses in each period the worst-case consumer information after observing

[3]See also: Haghpanah and Siegel (2020), Haghpanah and Siegel (forthcoming), Deb and Roesler (2021) and Yang (2021) for multi-product counterparts; as well as Doval and Skreta (2021) for a monopolist with limited commitment.

the realized price. Yang (2019) introduces an informational intermediary who contracts with the seller (with a private cost) on what information to disclose to consumers. The seller then chooses an optimal selling mechanism based on the information structure provided by the intermediary. In this paper, we consider a dynamic game where each stage game involves a monopolist posting a price first, and then an intermediary choosing an information structure for consumers. From this regard, pricing and information disclosure occur in every period in our model.[4] Moreover, since the informational intermediary in our paper is long-lived, future continuation plays could affect current information disclosed by the intermediary, which in turn affects outcomes.

Meanwhile, since the intermediary seeks to enhance and maintain its prominence, which in turn depends on the intermediary's behavior in the past, our paper is also related to the reputation literature, including its general theory (e.g., Fundenberg and Levine (1989) and Fundenberg and Levine (1992)); its effects on firms' competition and on inducing efficient effort levels (e.g., Mailath and Samuelson (2001) and Hörner (2002)); and its effect on an expert's credibility and their ability to communicate information (e.g., Ely and Välimäki (2003), Ottaviani and Sørensen (2006), and Vong (2021b)). A key distinction of our paper is that we abstract from the endogenous formulation of reputation and model the intermediary's concern about its prominence via an exogenous market-evolution process, so that the market growth rate in each period is an (exogenous) function of consumer surplus in the same period, whose shape can in turn be interpreted as results of different communication network among consumers and their word-of-mouth behaviors (see, for example, Chakraborty, Deb, and Öry (2021)); different rating systems (see, for example, Che and Hörner (2018), Hörner and Lambert (2021), and Vellodi (2021)); or the degree of competition (see Section 7). In the meantime, our assumption that the market growth rate depends only on consumer surplus in the same period resembles models where there is only limited record (e.g., Liu (2011) and Liu and Skrzypacz (2014)). In particular, while obtained by different reasons, our unintended welfare loss result has a similar flavor as the result of Bar-Isaac and Deb (2021), who show that more audiences observing past histories in the future may lead to worse outcomes for an agent with reputation concerns.

[4] Furthermore, even when restricting to stationary-Markov perfect equilibria, a stage game in our model is different from the aforementioned papers: The seller chooses selling mechanisms either after or while the information structure is chosen in Roesler and Szentes (2017) and Ravid, Roseler, and Szentes (forthcoming), respectively; while the seller chooses a price before an information structure is chosen in this paper. The seller can commit to any selling mechanism and the Nature always plays against the seller in Du (2018) and in Libgober and Mu (2021); while the seller is restricted to posted price mechanisms and the information structure is chosen to maximize a linear combination of the sales revenue and the consumer surplus in this paper. In the meantime, our model is similar to that of Libgober and Mu (2021) in that information structures can be contingent on the posted prices.

Methodologically, this paper is related to the Bayesian persuasion literature.[5] Specifically, the intermediary's disclosure problem in our model can be regarded as a persuasion problem when only the expected value of the state is payoff relevant (see Gentzkow and Kamenica (2016) and Dworczak and Martini (2019)). Meanwhile, the underlying structure of the game can be regarded as a dynamic game with a long-lived player and a short lived player, as studied by Fundenberg and Levine (1989) and Fudenberg, Kreps, and Maskin (1990), with an exception that there is a history-dependent state that scales the stage game payoffs. Moreover, the equilibrium selection criterion we use for this dynamic game in Section 4 follows the spirit of Maskin and Tirole (2001).

In terms of applications, this paper is also related to the certification literature (e.g., Biglaiser (1993), Lizzeri (1999), Stahl and Strausz (2017), Harbaugh and Rasmusen (2018), Vong (2021a), and Ali, Haghpanah, Lin, and Siegel (2022)) and the recent literature on online influencers. In the certification literature, our model is closer to that of Lizzeri (1999), who studies the optimal disclosure policy of a certifier in a market featuring adverse selection who can charge the seller a fee in exchange of (credibly) disclosing some information about the product to buyers. While our informational intermediary also discloses product information to consumers, she does so after seeing the seller's price. Moreover, there is no adverse selection problem in our setting without the presence of the intermediary as sellers do not possess any private information.

In the literature of online influencers, Fainmesser and Galeotti (2021) study a market where influencers can be paid to endorse products as sponsored recommendation, but with an opportunity cost of recommending fewer carefully-selected high-quality products to consumers, which would in turn affect their follower bases. Our model shares a similar spirit in which the intermediary faces a trade-off between prominence and short-term revenue. Nonetheless, in addition to having a dynamic model instead of a static one, we focus on the information provision aspect of an intermediary's service, rather than product recommendation and endorsement. Moreover, we highlight the business model where intermediaries are paid through commission rather than a lump-sum transfer. Mitchell (2021) also examines the economic implications of an influencer's trade-off between advertisement contents and good advice, and uses techniques in the dynamic contracting literature to characterize the optimal dynamic contracts for the follower. Our paper is complementary in the sense that we abstract from consumers' long-term endogenous relationships with the intermediary and focus on the details of pricing and information provision, while Mitchell (2021) abstracts from pricing and information provision and examines the relationship between the follower and the influencer. Pei and Mayzlin (forthcoming) studies influencers' paid promotion contents using a Bayesian persuasion framework. In their model, a product seller can pay an influencer to increase the

[5] See Kamenica (2018) for a comprehensive review.

likelihood of a positive signal for a given two-state-two-signal Blackwekk experiment. The main difference between their model and ours is that our intermediary can use any Blackwell experiment to inform consumers about their values, which cannot be altered by the seller ex-post.

3 Model

3.1 Primitives

Time $t \geq 0$ is discrete. In each period t, a short-lived seller sells one product to a mass $m_t \geq 0$ of short-lived consumers with unit demands through a long-lived intermediary with discount factor $\delta \in (0,1)$. A consumer in period t is denoted by $(v,x) \in \mathbb{R}_+ \times [0,m_t]$, where v denotes the consumer's value for the product and x denotes consumers' payoff-irrelevant identities. Across consumers, v is distributed according to a probability measure described by a decreasing function $\overline{D} : \mathbb{R}_+ \to [0,1]$, so that $\overline{D}(p)$ is the share of consumers with $v \geq p$; whereas x is distributed according to the Lebesgue measure, so that the mass of consumers in period t is m_t. Assume that $\overline{D}$ is *regular*, in the sense that it is continuously differentiable, has a non-zero derivative on an interval in $\mathbb{R}_+$, and induces a decreasing marginal revenue function.[6] Each consumer knows their payoff-irrelevant type x but has to learn about the product value v through the information provided by the intermediary.

3.2 Timing and Payoffs

In each period, the timing of moves is as follows: (i) Consumers arrive, (ii) the seller posts a price, (iii) the intermediary observes the posted price and then provides information about v to consumers (see Section 3.3), (iv) consumers then decide whether to buy the product after receiving information and observing the posted price p, (v) payoffs are realized and consumers observe their values v. Given the outcome, a consumer has payoff $v - p$ if he buys the product and has payoff zero if he does not buy, while the seller and the intermediary share the revenue by a fixed proportion $\alpha \in (0,1)$. That is, the seller's payoff is $1-\alpha$ share of the total revenue and the intermediary's (stage-game) payoff is α share of the total revenue, where the total revenue is the posted prices times the trade volume.

The mass of consumers m_t who trade through the intermediary (i.e., the market size) in period t depends on outcomes in previous periods. In period 0, the market size is 1 (i.e., $m_0 = 1$). In each period $t \geq 0$, the market size m_{t+1} in period $t+1$ is the integral of an increasing affine function of ex-post surplus across all consumers in period t. In other words,

[6] That is, the function $q \mapsto q\overline{D}^{-1}(q)$ is concave. This is equivalent to assuming that $1 - \overline{D}$ is regular in the Myersoninan sense.

after making the purchasing decision at posted price p in period t, the market size m_{t+1} of period $t+1$ is given by:

$$m_{t+1} = \max\left\{\int_0^{m_t}\int_0^{\infty}[\gamma + \beta \cdot a(p,v,x)(v-p)]\overline{D}(\mathrm{d}v)\,\mathrm{d}x, 0\right\}, \tag{1}$$

for some $\beta \geq 0$ and $\gamma \in [0, 1/\delta)$, where $a(p,v,x)$ is the probability that consumer (v,x) purchases the product at price p, and β measures the *market feedback* as it governs how much past consumers surplus could affect future market sizes.

3.3 Information

In each period, consumers receive information about v from the intermediary. A *disclosure policy* is a measurable function $\phi : \mathbb{R}_+ \times [0,1] \to \mathbb{R}$. Given any market size $m > 0$ and any disclosure policy ϕ, consumer $x \in [0,m]$ observes $\phi(v, x/m)$. In other words, a disclosure policy is an individualized partition that discloses information about v to consumers by revealing them which element of the partition does the product value belong to. The intermediary can choose any disclosure policy when providing information to the consumers.

Given any market size $m > 0$ and any disclosure policy ϕ, consumer $x \in [0,m]$ has interim expected value[7]

$$\mathbb{E}\left[v\middle|\phi\left(v, \frac{x}{m}\right)\right].$$

For any $p \geq 0$, let

$$\begin{aligned} D_\phi(p) :=& \frac{1}{m}\int_0^m\int_0^\infty \mathbf{1}\left\{\mathbb{E}\left[v\middle|\phi\left(\tilde{v}, \frac{x}{m}\right)\right] \geq p\right\}\overline{D}(\mathrm{d}\tilde{v})\,\mathrm{d}x \\ =& \int_0^1\int_0^\infty \mathbf{1}\{\mathbb{E}[v|\phi(\tilde{v},u)] \geq p\}\overline{D}(\mathrm{d}\tilde{v})\,\mathrm{d}u \end{aligned}$$

denote the (normalized) marginal distribution of consumers' interim expected values. By the law of iterated expectation, for any market size m and for any disclosure policy ϕ,

$$\mathbb{E}\left[\mathbb{E}\left[v\middle|\phi\left(v, \frac{x}{m}\right)\right]\right] = \mathbb{E}[v].$$

Thus, D_ϕ must be a mean-preserving contraction of $\overline{D}$. In fact, the converse is also true (see Green and Stokey (1978) and Gentzkow and Kamenica (2017)): For any distribution D that is a mean-preserving contraction of $\overline{D}$, there exists a disclosure policy ϕ such that $D_\phi = D$.

As a result, we may represent the set of all possible disclosure policies by the collection of distributions that are mean-preserving contractions of $\overline{D}$. That is, a disclosure policy is

[7] Notice that x/m can be regarded as a random variable that is uniformly distributed on $[0,1]$ and is independent of v. Thus, the conditional expectation is well-defined.

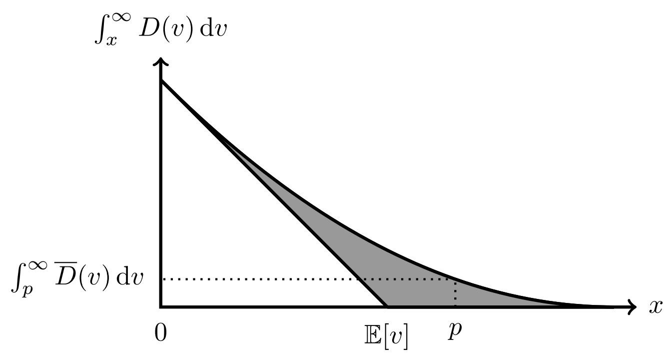

Figure 1: Feasible Disclosure Policies $\mathcal{D}$

$D \in \mathcal{D}$, where $\mathcal{D}$ is the collection of nonincreasing, upper-semicontinuous functions $D : \mathbb{R}_+ \to [0,1]$ such that $D(0) = 1$ and

$$\int_p^\infty D(v)\,\mathrm{d}v \leq \int_p^\infty \overline{D}(v)\,\mathrm{d}v, \tag{2}$$

for all $p \geq 0$, with equality at $p = 0$. The set $\mathcal{D}$ is described in Figure 1, where each $D \in \mathcal{D}$ corresponds to a decreasing and convex function whose graph is in the highlighted area.

Using this representation, the model can be simplified. First, notice that since consumers are short-lived and make purchasing decisions after observing the price p_t, the disclosure policy ϕ_t and the value of ϕ_t, they must purchase the product if their interim expected value is above p_t, and not purchase the product if the interim expected value is below p_t. Moreover, since $D \in \mathcal{D}$ corresponds to the marginal distribution of consumers' interim expected values under disclosure policy D, the trade volume would be $m \cdot q$ for some $q \in [D(p^+), D(p)]$ when the market size is m, the disclosure policy is $D \in \mathcal{D}$, and the posted price is p (i.e., $m \cdot D(p)$ is the *demand* at price p under disclosure policy D). This in turn implies that the seller's share of revenue would be $(1-\alpha)m \cdot p \cdot q$ and the intermediary's share of revenue would be $\alpha m \cdot p \cdot q$, for some $q \in [D(p^+), D(p)]$, and the exact q depends on consumers' tie-breaking rule.

Moreover, the law of motion (1) for the market sizes can also be simplified. To see this, notice that given any market size $m_t > 0$, any posted price p_t and any disclosure policy ϕ_t, consumers' optimal purchasing decision $a(p_t, v, x) \in [0,1]$ given (v, x) and p_t must be such that

$$a(p,v,x) = \begin{cases} 1, & \text{if } \mathbb{E}\left[v \,\middle|\, \phi\left(v, \frac{x}{m_t}\right)\right] > p_t \\ 0, & \text{if } \mathbb{E}\left[v \,\middle|\, \phi\left(v, \frac{x}{m_t}\right)\right] < p_t \end{cases}$$

Thus, by the definition of D_{ϕ_t} and by the law of iterated expectation, (1) can be written as

$$m_{t+1} = m_t\left(\gamma + \beta\int_0^\infty \mathbf{1}\{v \geq p_t\}(v - p_t)D_{\phi_t}(\mathrm{d}v)\right) = m_t\left(\gamma + \beta\int_{p_t}^\infty D_{\phi_t}(v)\,\mathrm{d}v\right).$$

Therefore, since the set of disclosure policies is represented by $\mathcal{D}$, given any market size $m_t \geq 0$, any posted price $p_t \geq 0$, and any disclosure policy $D_t \in \mathcal{D}$, (1) can be simplifed to

$$m_{t+1} = m_t\left(\gamma + \beta\int_{p_t}^\infty D_t(v)\,\mathrm{d}v\right). \tag{3}$$

3.4 Strategies and Solution Concepts

Since the consumers' behavior can be simplified and embedded into the choice of $D \in \mathcal{D}$ up to a tie-breaking rule (as noted above), the model can effectively be reduced to a perfect information extensive form game with a long-lived intermediary, a (sequence of) short-lived seller, and a (sequence of) short-lived tie-breaker whose payoff is constant across all histories. In each period t, the market size m_t is determined by (3). The seller moves first and posts a price $p_t \geq 0$. The intermediary then observes the price and chooses a disclosure policy $D_t \in \mathcal{D}$. The tie-breaker observes both p_t and D_t and chooses $q_t \in [D_t(p_t^+), D(p_t)]$. Payoffs are then realized so that the intermediary's and the seller's payoff are $\alpha m_t \cdot p_t \cdot q_t$ and $(1-\alpha)m_t \cdot p_t \cdot q_t$, respectively.

As a result, our model can be viewed as a dymamic game with one long-run player (i.e., the intermediary) and two (sequences of) short-run players (i.e, a seller and a tie-breaker in each period) (see Fudenberg, Kreps, and Maskin (1990)), with two exceptions: (i) the moves are sequential in each stage game,[8] and (ii) there is a deterministically evolving payoff-relevant state variable (i.e., the market size).[9] Henceforth, we represent our model using this dynamic game. In this game, the seller's strategy is a mapping from past histories[10] to $\mathbb{R}_+$, the intermediary's strategy is a mapping from past histories and the seller's price to $\mathcal{D}$, and the tie-breaker's strategy is a mapping from past histories, the seller's price p_t, and the intermediary's disclosure policy D_t to $[D_t(p_t^+), D_t(p_t)]$.

Notice that although stage game payoffs depend on the market size, they are simply scaled. In particular, market sizes do not affect the players' preference over stage game outcomes. In the spirit of Maskin and Tirole (2001), we say that a strategy profile is *stationary-Markov* if it does not depend on past histories. That is, a strategy profile is stationary-Markov if

[8] This means that when checking the intermediary's deviations in a period t, we need to consider all histories induced by prices charged by the seller in period t, even if only one of them would be on-path.

[9] One implication of this is that, if the market size is not bounded, the intermediary's payoff could be unbounded. As a result, there might be subgame perfect equilibria in which the intermediary's payoff diverges at some histories.

[10] See the online appendix for a formal definition of histories

the seller's price does not depends on past histories, the intermediary's disclosure policy only depends on the seller's price in the same period, and the tie-breaker's strategy only depends on the seller's price and the intermediary's disclosure policy in the same period. A subgame perfect equilibrium is said to be *stationary-Markov* (or simply *Markov* hereafter) if it is a stationary-Markov strategy profile. Moreover, a subgame perfect equilibrium is said to be *finite* if the intermediary's continuation value at any history is finite, and is *infinite* if otherwise. In Section 4, we first characterize the set of Markov perfect equilibrium outcomes. Then, Section 5 further characterizes the entire feasible payoff set among all subgame perfect equilibria.

3.5 Discussions of the Model

Short-lived seller—The seller is assumed to be short-lived. This captures a feature in markets of informational intermediaries where the collaborating relationships between sellers and intermediaries are often not long-term. Rather, sellers often search for intermediaries to collaborate with whenever they introduce a new product. This feature is one of the driving forces of our main result, as it creates a difference in interest between the seller and the intermediary. In Section 6.3, we consider an alternative setting where the seller is long-lived and has the same discount rate as the intermediary's, capturing a benchmark where the seller and the intermediary are able to form long-term collaborating relationships. In turns out that the unintended welfare loss would be eliminated when the seller is long-lived, which highlights an undesirable welfare consequence of short-term collaborating relationships between sellers and intermediaries.

Market structure—We assume that both the product seller and the intermediary are monopolists. An interpretation is that while there are many sellers and many intermediaries in the market, each of them has some market powers. When collaborating with a seller, the intermediary provides both an exclusive access to a group of consumers (whose volume depends on its prominence) and information. Together with the short-lived seller assumption, our model can be regarded as the longitudinal section of an intermediary's life cycle.

Disclosure policy—A disclosure policy in our setting is equivalent to a Blackwell experiment of a single buyer's value. Therefore, our model of information follows the literature of consumer information and pricing, where consumers are ex-ante uninformed and can receive any signal for their own values. As motivated in this literature (see, for instances, Johnson and Myatt (2006) and Anderson and Renault (2006)), these signals can be interpreted as information about characteristics of a product. The more information consumers have, the better they know about their (private) match values for this product. In other words, we focus on the aspect where the intermediary provides "horizontal information" rather than "vertical information" such as product quality; while consumers choose between buying or not, rather than which product to buy. Consequently, our model abstracts from *paid sponsorship*, a business model sometimes used by intermediaries, especially influencers, where sellers pay intermediaries for them to promote their products. Instead, we only consider the information-provision service of intermediaries.

Revenue sharing rule—The seller and the intermediary are assumed to share total sales revenue in each period through an exogenous and constant sharing rule $\alpha \in (0, 1)$. This reflects a feature of the market of informational intermediary that it is sometimes difficult to contingent prices of intermediaries' services on their prominence, due to either legal, technical, or institutional reasons.[11] In Section 7, we extend our model and show that there would still be unintended welfare losses even if the sharing rule (under parametric assumptions) depends on the intermediary's prominence level

Market evolution—The market sizes $\{m_t\}$ is assumed to evolve through an exogenously given law of motion, with a growth rate that depends linearly on the consumer surplus in the same period. The linearity assumption is mostly for the ease of exposition, as it admits a closed-form solution for the intermediary's problem when taking her continuation value as given. In Section 7, we relax this assumption and consider more general mappings from consumer surplus to growth rates. In contrast, the exogenous law of motion is a more consequential assumption, as it abstracts from strategic decisions of consumers in terms of entry and exit. In Section 7, we introduce a specific micro foundation for this exogenous law of motion, where consumers are short-lived but some of them can decide whether they would exit and obtain their (heterogeneous) outside options.

Effects of higher market feedback—Under our assumptions, higher market feedback level β has three effects: (i) it leads to a more market-minded intermediary who assigns more weights to consumer surplus; (ii) it leads to higher market growth rate; and (iii) it may change the equilibrium strategies. It is noteworthy that effect (i) is always (weakly) beneficial for consumers, while effect (ii) is always (weakly) beneficial for every players. Nonetheless, our main result suggests that higher β may lead to Pareto worse outcomes despite the positive effects of (i) and (ii). In other words, the essence of our main results is effect (iii) could possibly dominate (i) and (ii).

[11] For example, it would be difficult for product sellers to contingent their commission rates for financial and insurance brokers, since their prominence level may be difficult to measure. Alternatively, influencers retain commissions through third party affiliate links (e.g., rStyle or ShopStyle) and hence the commission rates for influencers mostly remain the same for similar products regardless of how many followers an influencer has.

4 Markov Perfect Equilibrium and Unintended Welfare Loss

In this section, we characterize all Markov perfect equilibrium outcomes and show that higher market feedback might lead to Pareto-inferior outcomes. To this end, we first establish formally the criterion we use for such comparisons. Given any strategy profile, we refer to the sequence of normalized sales revenues (i.e., sales revenue divided by the market size), normalized consumer surpluses, the intermediary's normalized continuation payoffs, prices, and market sizes in each period as an *outcome*, which we denote by $\mathbf{z} := \{r_t, \sigma_t, \omega_t, p_t, m_t\}$.

Definition 1. For any two outcomes $\mathbf{z} = \{r_t, \sigma_t, \omega_t, p_t, m_t\}$ and $\mathbf{z}' = \{r'_t, \sigma'_t, \omega'_t, p'_t, m'_t\}$, say that $\mathbf{z}'$ is *dominated* by $\mathbf{z}$ (denoted as $\mathbf{z} \succ \mathbf{z}'$) if for all t, $m'_t r'_t \leq m_t r_t$, $m'_t \sigma'_t \leq m_t \sigma_t$ and $m'_t \omega'_t \leq m_t \omega_t$, with at least one inequality being strict.

Moreover, when restricting to Markov perfect equilibria, as players' strategies do not depend on histories in previous periods, in any Markov perfect equilibrium outcome, only the market sizes would depend on t. As a result, we may suppress the unnecessary subscripts and write a Markov perfect equilibrium outcome as $\mathbf{z}^{\mathrm{M}} = (r^{\mathrm{M}}, \sigma^{\mathrm{M}}, \omega^{\mathrm{M}}, p^{\mathrm{M}}, \{m_t^{\mathrm{M}}\})$.

4.1 Finite Markov Perfect Equilibrium Outcomes

In what follows, we first focus on finite Markov perfect equilibrium outcomes. We characterize all finite Markov perfect equilibrium outcomes and compare them by varying market feedback β. Let $\overline{p}$ be the unique solution of $\max_p p\overline{D}(p)$, and let

$$\underline{\beta} := \frac{1-\gamma\delta}{\delta\mathbb{E}[v]} \quad \text{and} \quad \overline{\beta} := \frac{1-\gamma\delta}{\delta\int_{\overline{p}}^{\infty}\overline{D}(v)\,\mathrm{d}v}.$$

Below we first state our main welfare result.

Proposition 1 (Unintended Welfare Loss—Finite Markov)**.** *For any $\beta < \overline{\beta}$, there exists a unique finite Markov perfect equilibrium outcome $\mathbf{z}^{\mathrm{M}}(\beta)$. Furthermore, $\mathbf{z}^{\mathrm{M}}(\beta) \succ \mathbf{z}^{\mathrm{M}}(\beta')$ for all β, β' such that $\underline{\beta} < \beta < \beta' < \overline{\beta}$.*

According to Proposition 1, higher market feedback does not necessarily benefit the consumers, even if it compels the intermediary to be concerned more about consumer surplus. In fact, higher market feedback might even lead to Pareto inferior outcomes. Whenever $\beta \in (\underline{\beta}, \overline{\beta})$, an increase in market feedback would lead to unintended welfare loss by yielding Pareto inferior finite Markov perfect equilibrium outcomes.

Proposition 1 is established by characterizing all the finite Markov perfect equilibrium outcomes. We begin the analysis by noting that even though payoffs could potentially be unbounded, the one-shot deviation principle would still hold when considering strategy profiles that yield a finite continuation payoff for the intermediary at every history. This is because stage game payoffs are bounded from below and because the intermediary's payoff is additively separable. Lemma 1 below summarizes this observation.

Lemma 1 (One-Shot Deviation Principle)**.** *Given any strategies of the seller and the tie-breaker, for any history h^t in any period t, and for any strategy of the intermediary that yields a finite continuation payoff, there is a profitable deviation from the continuation strategy at h^t if and only if there is a profitable one-shot deviation at some history after h^t.*

We now outline the characterization of finite Markov perfect equilibrium outcomes. To begin with, notice that with Lemma 1, Markov perfect equilibria can be characterized by the incentives of both the intermediary and the (short-lived) seller in each period while holding each other's strategy fixed. This leads to the following lemma.

Lemma 2. *A finite Markov perfect equilibrium is characterized by a tuple $(\omega^{\mathrm{M}}, p^{\mathrm{M}}, \mathbf{D}^{\mathrm{M}})$ with $\omega^{\mathrm{M}}, p^{\mathrm{M}} \in [0, \infty)$ and $\mathbf{D}^{\mathrm{M}} : \mathbb{R}_+ \to \mathcal{D}$ that satisfy the following conditions:*

$$\omega^{\mathrm{M}} = \sup_{D \in \mathcal{D}} \left[\alpha p^{\mathrm{M}} D(p^{\mathrm{M}}) + \delta \omega^{\mathrm{M}} \left(\gamma + \beta \int_{p^{\mathrm{M}}}^{\infty} D(v) \, \mathrm{d}v \right) \right], \tag{4}$$

$$p^{\mathrm{M}} \mathbf{D}^{\mathrm{M}}(p^{\mathrm{M}} | p^{\mathrm{M}}) \geq p \mathbf{D}^{\mathrm{M}}(p|p), \tag{5}$$

for all $p \geq 0$,

$$\alpha p \mathbf{D}^{\mathrm{M}}(p|p) + \delta \omega^{\mathrm{M}} \left(\gamma + \beta \int_{p}^{\infty} \mathbf{D}^{\mathrm{M}}(v|p) \, \mathrm{d}v \right) \geq \alpha p D(p) + \delta \omega^{\mathrm{M}} \left(\gamma + \beta \int_{p}^{\infty} D(v) \, \mathrm{d}v \right), \tag{6}$$

for all $p \geq 0$ and for all $D \in \mathcal{D}$. Furthermore, for any Markov perfect equilibrium $(\omega^{\mathrm{M}}, p^{\mathrm{M}}, \mathbf{D}^{\mathrm{M}})$, its outcome is given by $(r^{\mathrm{M}}, p^{\mathrm{M}}, \sigma^{\mathrm{M}}, \omega^{\mathrm{M}}, \{m_t^{\mathrm{M}}\})$, where $r^{\mathrm{M}} = p^{\mathrm{M}} \mathbf{D}^{\mathrm{M}}(p^{\mathrm{M}} | p^{\mathrm{M}})$, $\sigma^{\mathrm{M}} = \int_{p^{\mathrm{M}}}^{\infty} \mathbf{D}^{\mathrm{M}}(v|p^{\mathrm{M}}) \, \mathrm{d}v$, and $m_t^{\mathrm{M}} = (\gamma + \beta \sigma^{\mathrm{M}})^t$ for all $t \geq 0$.

One of the implications of this characterization is that the tie-breaker always breaks tie to maximize sales revenue in each period, and thus it would not be necessary to keep track of the tie-breaker's strategy when restricting to finite Markov perfect equilibria.

Using Lemma 2, characterizing finite Markov perfect equilibria becomes essentially a static problem that consists of: (i) solving for the intermediary's per-period best response, given a price p posted by the seller in that period and given its continuation value (solving (6) given p, ω^{M}); (ii) solving for the seller's best response given (i) (solving (5) given ω^{M} and the intermediary's best response derived in (i)); and (iii) finding a consistent continuation payoff (verifying (4) given (i) and (ii)). We now introduce two lemmas that characterize the solutions of steps (i) and (ii) for a fixed continuation value.

For any $p \geq 0$, let $v(p) := \mathbb{E}[v | v \geq p]$ and let $v^{-1}(p) := \inf\{x \geq 0 | v(x) \geq p\}$.[12] Notice that both v and v^{-1} are nondecreasing and $v^{-1}(p) = 0$ for all $p \in [0, \mathbb{E}[v]]$.

[12] As a convention, if p is greater than the upper bound of the support of $\overline{D}$, define $\mathbb{E}[v|v \geq p]$ as this upper bound. Meanwhile, if p is greater than $\max_{x \geq 0} v(x)$, then define $v^{-1}(p)$ as $\max_{x \geq 0} v(x)$.

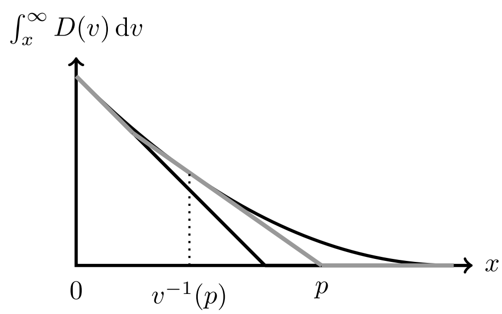

(a) Optimal cutoff $= v^{-1}(p)$

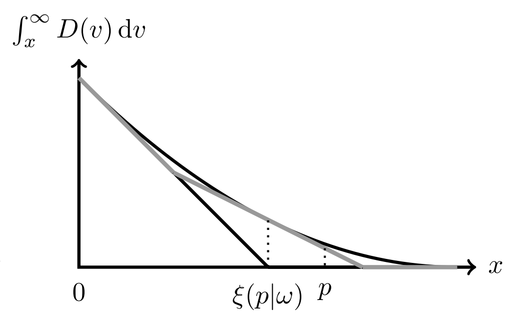

(b) Optimal Cutoff $= \xi(p|\omega)$

Lemma 3. *For any $p, \omega \in [0, \infty)$,*

$$\Delta(p|\omega) := \underset{D \in \mathcal{D}}{\operatorname{argmax}} \left[\alpha p D(p) + \delta\omega \left(\gamma + \beta \int_p^\infty D(v)\,\mathrm{d}v \right) \right]$$

is nonempty. Moreover, for any $D \in \Delta(p|\omega)$, $D(v) = \overline{D}(\xi(p|\omega))$, for all $v \in [\xi(p|\omega), p]$ and

$$\int_{\xi(p|\omega)}^\infty D(v)\,\mathrm{d}v = \int_{\xi(p|\omega)}^\infty \overline{D}(v)\,\mathrm{d}v,$$

where

$$\xi(p|\omega) := \max\left\{ \left(1 - \frac{\alpha}{\delta\beta\omega}\right)^+ p, v^{-1}(p) \right\}.$$

Lemma 3 provides a characterization of the intermediary's optimal disclosure policy given a price p and a continuation value ω. For any $p, \omega \in [0, \infty)$, the intermediary essentially faces a static problem where she chooses a demand $D \in \mathcal{D}$ to maximizes a linear combination of the sales revenue and consumer surplus.

To gain some intuitions, consider first the case when $\beta = 0$ so that the intermediary seeks to maximize sales revenue, which is equivalent to maximizing sales volume for each given price. In this case, for any price $p \in [0, \mathbb{E}[v]]$, the intermediary can simply disclose no information and every consumer would buy. Meanwhile, if $p \in (\mathbb{E}[v], \infty)$, then the largest possible trade volume can be attained if every consumer with values above a threshold buys and if their interim expected value equals to p, which in turn implies that the threshold is $v^{-1}(p)$ (see Figure 2a). Notice that in both cases, consumer surplus equals to zero. Now suppose that $\beta > 0$, then the intermediary would benefit from leaving the consumers some surplus. This means that she may rather induce fewer consumers to buy in order to prevent low-value consumers from buying at a high price. As a result, the intermediary would prefer inducing a higher threshold, which is precisely given by $\xi(p|\beta)$ (see Figure 2b).

For any $\omega \in [0, \infty)$, anticipating the intermediary's best response, the seller effectively solves a revenue maximization with the demand at price p being the sales volume induced by the intermediary's best response against price p. Lemma 4 below characterizes the solution of this problem.

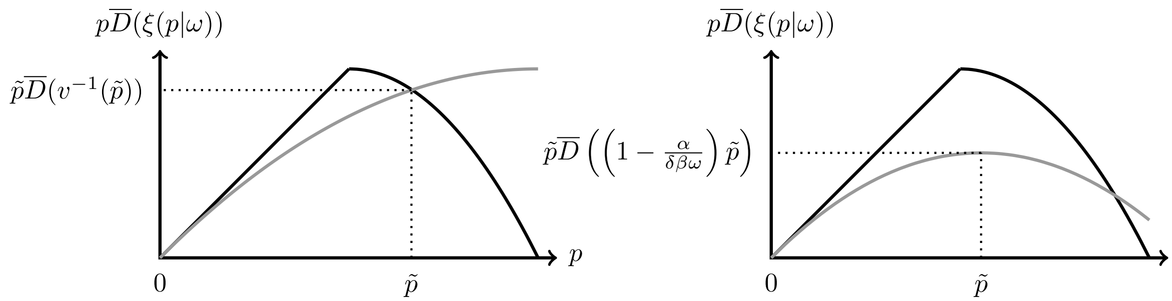

(a) First Inequality of (7) Binds (b) Second Inequality of (7) Binds

Figure 3: Optimal Price $\tilde{p}$

Lemma 4. *For any $\omega \in [0,\infty)$ and for any selection $\mathbf{D}$ of $\Delta(\cdot|\omega)$, the maximization problem*

$$\max_{p\geq 0} p\mathbf{D}(p|p)$$

has a unique solution $\tilde{p}$. Furthermore,

$$v^{-1}(\tilde{p}) \leq \left(1-\frac{\alpha}{\delta\beta\omega}\right)^{+}\tilde{p} \leq \overline{p}, \tag{7}$$

with at least one inequality binding. In particular,

$$\int_{\tilde{p}}^{\infty} \mathbf{D}(v|\tilde{p})\,\mathrm{d}v = 0 \iff \left(1-\frac{\alpha}{\delta\beta\omega}\right)^{+}\tilde{p} = v^{-1}(\tilde{p}).$$

To better understand Lemma 4, notice that by Lemma 3, $p\mathbf{D}(p|p) = p\overline{D}(\xi(p|\omega))$ for any $p,\omega \in [0,\infty)$ and for any selection $\mathbf{D}$ of $\Delta(\cdot|\omega)$. As a result, the seller's revenue maximization problem can be written as

$$\max_{p\geq 0} p\overline{D}(\xi(p|\omega)) = \max_{p\geq 0}\left[\min\left\{p\overline{D}\left(\left(1-\frac{\alpha}{\delta\beta\omega}p\right)^{+}\right), p\overline{D}(v^{-1}(p))\right\}\right].$$

Clearly, if $\delta\beta\omega \leq \alpha$, then the function above would coincide with $p\overline{D}(v^{-1}(p))$ and the optimal price must be $\mathbb{E}[v]$ and the first inequality of (7) binds and consumer surplus is zero. Meanwhile, if $\delta\beta\omega > \alpha$, then there are two possibilities, as depicted by Figure 3. The first possibility is as illustrated by Figure 3a, where the optimal price is the price at which the graph of the two functions intersect. In this case, the first inequality of (7) binds and the consumer surplus is zero. Another possibility is illustrated by Figure 3b. In this case, the optimal price is at which the first function is maximized. In this case, the second inequality of (7) binds and the consumer surplus is positive.

Even with a fixed ω, Lemma 4 already highlights the main driving force behind Proposition 1. Indeed, when β is close enough to zero (i.e., when $\delta\beta\omega \leq \alpha$), the induced sales revenue is $\mathbb{E}[v]$ and the allocation is efficient. As β increases, the optimal price would first increase while the consumer surplus would remain zero. This means that the sales revenue would decrease, consumers would not be benefited and the market sizes would remain the same as the case of lower β, leading to a Pareto-inferior outcome. Only when β is large enough will consumer surplus become positive and will the price begin to decrease. The intuition behind this feature is reminiscent of the hold-up problem logic: Since the intermediary discloses information after observing the seller's price, and since the intermediary seeks to enhance consumer surplus in addition to the sales revenue, the seller—in the anticipation of the intermediary's response—would charge a higher price to suppress the intermediary's urge of improving consumer surplus so that her interest could be more well-aligned with the seller's. This in turn leads to Pareto inferior outcomes

The above reasoning holds for any given continuation value ω. To complete proof of Proposition 1, we need to characterize the equilibrium continuation value. With Lemma 3 and Lemma 4, finite Markov perfect equilibrium outcomes can be completely characterized. Indeed, given Lemma 3 and Lemma 4 characterizes the solution of the seller's revenue maximization problem given ω, it remains to find the proper continuation value ω^{M} and equilibrium price p^{M} that satisfies (4), (5) and (6). To describe finite Markov perfect equilibrium outcomes, it is convenient to define

$$g^{\beta}(p) := \frac{\alpha}{\delta\beta}\left(1 + \frac{p\overline{D}(p)}{\int_p^{\infty}\overline{D}(v)\,\mathrm{d}v}\right),$$

for all $p \in [0, \overline{p}]$. Notice that the function $(\beta, p) \mapsto g^{\beta}(p)$ is continuous, strictly decreasing in β, and strictly increasing in p on $[0, \overline{p}]$. Meanwhile, let

$$p^{\beta} := \inf\left\{p \geq 0 \,\middle|\, \delta\left(\gamma + \beta\int_p^{\infty}\overline{D}(v)\,\mathrm{d}v\right) \geq 1\right\}.$$

By definition, $p^{\beta} \in (0, \overline{p})$ whenever $\beta \in (\underline{\beta}, \overline{\beta})$. Moreover, p^{β} is strictly decreasing in β on $[\underline{\beta}, \overline{\beta}]$. With this notation, finite Markov perfect equilibrium outcomes can be characterized by Theorem 1 below.

Theorem 1 (Finite Markov Perfect Equilibrium Outcomes)**.** *A finite Markov perfect equilibrium outcome exists if and only if $\beta \in [0, \overline{\beta}]$. Furthermore, the following are equivalent:*

1. $\mathbf{z}^{\mathrm{M}} = (r^{\mathrm{M}}, \sigma^{\mathrm{M}}, \omega^{\mathrm{M}}, p^{\mathrm{M}}, \{m_t^{\mathrm{M}}\})$ *is a finite Markov perfect equilibrium outcome.*
2. $\omega^{\mathrm{M}} \geq g^{\beta}(\overline{p})$ *if* $\beta = \overline{\beta}$, *while*

$$\omega^{\mathrm{M}} = \begin{cases} \frac{\alpha\mathbb{E}[v]}{1-\gamma\delta}, & \text{if } \beta \in [0, \underline{\beta}] \\ g^{\beta}(p^{\beta}), & \text{if } \beta \in (\underline{\beta}, \overline{\beta}) \end{cases}.$$

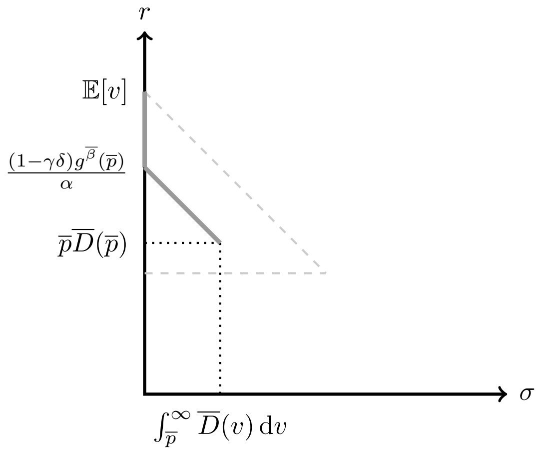

Figure 4: Surplus Divisions under Finite Markov Perfect Equilibria

Moreover,

$$p^{\mathrm{M}} = \begin{cases} \mathbb{E}[v], & \text{if } \beta \in [0, \underline{\beta}] \\ v(p^{\beta}), & \text{if } \beta \in (\underline{\beta}, \overline{\beta}) \\ \frac{\delta\beta\omega^{\mathrm{M}}}{\delta\beta\omega^{\mathrm{M}}-\alpha}\overline{p}, & \text{if } \beta = \overline{\beta} \end{cases} ; \quad r^{\mathrm{M}} = \begin{cases} \mathbb{E}[v], & \text{if } \beta \in [0, \underline{\beta}] \\ \frac{(1-\gamma\delta)}{\alpha} g^{\beta}(p^{\beta}), & \text{if } \beta \in (\underline{\beta}, \overline{\beta}) \\ \frac{\delta\beta\omega^{\mathrm{M}}}{\delta\beta\omega^{\mathrm{M}}-\alpha}\overline{p}\overline{D}(\overline{p}), & \text{if } \beta = \overline{\beta} \end{cases} ;$$

$$\sigma^{\mathrm{M}} = \begin{cases} 0, & \text{if } \beta \in [0, \overline{\beta}) \\ \int_{\overline{p}}^{\infty} \overline{D}(v)\,\mathrm{d}v - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}-\alpha}\overline{p}\overline{D}(\overline{p}), & \text{if } \beta = \overline{\beta} \end{cases} ;$$

and $m_t^{\mathrm{M}} = (\gamma + \beta\sigma^{\mathrm{M}})^t$*, for all* $t \geq 1$*.*

With Theorem 1, Proposition 1 then follows immediately. Figure 4 plots the set of normalized surplus divisions induced by finite Markov perfect equilibria across all $\beta \in [0, \overline{\beta}]$. For any $\beta \in [0, \underline{\beta}]$, the sales revenue equals to $\mathbb{E}[v]$ and consumer surplus is zero. For any $\beta \in (\underline{\beta}, \overline{\beta})$, the sales revenue equals to $(1 - \gamma\delta)g^{\beta}(p^{\beta})/\alpha$ and consumer surplus remains zero. In particular, sales revenue decreases as β increases. Finally, when $\beta = \overline{\beta}$, there are multiple finite Markov perfect equilibrium outcomes and every such outcome induces a total surplus of $(1 - \gamma\delta)g^{\beta}(\overline{p})/\alpha$. Specifically, when $\beta \in (\underline{\beta}, \overline{\beta})$, higher market feedback leads to Pareto inferior outcomes. Moreover, even if consumer surplus can be positive when $\beta = \overline{\beta}$, the induced total surplus must lower than those induced by any $\beta < \overline{\beta}$. For comparison, Figure 4 also plots the set of possible surplus division under the same $\overline{D}$ that can arise under any consumer information (see Roesler and Szentes (2017)) when the seller prices after seeing the information structure, which corresponds to the area bounded by the dashed lines.

4.2 Infinite Markov Perfect Equilibria

The previous section restricts attention to finite Markov perfect equilibria outcomes. Compared with infinite Markov perfect equilibria, finite Markov perfect equilibria are arguably more selective and natural since the intermediary would be indifferent among many choices—especially off the equilibrium path—when the present discounted profit diverges, which in turn allows the intermediary to punish the seller's deviation severely. Nonetheless, understanding infinite Markov perfect equilibrium outcomes may further complete the characterization. In what follows, we further characterize all infinite Markov perfect equilibrium outcomes.

Since the market grows at a rate that is increasing in β for any fixed level of consumer surplus, the intermediary's present discounted profit may diverge in principle, even if there is discount. As noted above, when allowing for divergent payoffs, the intermediary would be indifferent among many different disclosure policies in a given period, as long as the continuation payoff diverges. This in turn allows the intermediary to punish the seller's deviation more severely, and thus there are generally many infinite Markov equilibrium outcomes. Nevertheless, even when considering infinite Markov perfect equilibria as well, the welfare implication of Proposition 1 still remains the same qualitatively.

To summarize this, we first let

$$r^* := \sup_{p \geq 0} \inf_{D \in \mathcal{D}} \inf_{q \in [D(p^+), D(p)]} p \cdot q \tag{8}$$

denote the *revenue guarantee*. That is, r^* is the sales revenue that the seller can secure in each period, regardless of the strategies of the intermediary and the tie-breaker. By definition, for any $r < r^*$, there exists $p \geq 0$ such that for any $D \in \mathcal{D}$, $pD(p^+) > r$.

Notice that for any $p \geq \mathbb{E}[v]$, $\inf_{D \in \mathcal{D}} pD(p^+) = 0$ since the demand $\underline{D} \in \mathcal{D}$ that has a jump of 1 at $\mathbb{E}[v]$ gives $p\underline{D}(p^+) = 0$. Meanwhile, for any $p \in [0, \mathbb{E}[v])$, there exists a unique $\zeta(p) \geq p$ such that $\mathbb{E}[v|v \leq \zeta(p)] = p$. It then follows that $\inf_{D \in \mathcal{D}} pD(p^+) = p\overline{D}(\zeta(p))$ (see Figure 5).[13] Therefore, we must have

$$r^* = \max_{p \in [0, \mathbb{E}[v]]} p\overline{D}(\zeta(p)).$$

Now let

$$\widetilde{\beta} := \frac{1 - \gamma\delta}{\delta(\mathbb{E}[v] - r^*)},$$

and note that $\underline{\beta} < \widetilde{\beta} < \overline{\beta}$. We then have the following.

Proposition 2 (Unintended Welfare Loss–Markov Perfect Equilibria)**.** *For any $\beta < \widetilde{\beta}$, there exists a unique Markov perfect equilibrium outcome $\mathbf{z}^{\mathrm{M}}(\beta)$ and any Markov perfect equilibrium that induces outcome $\mathbf{z}^{\mathrm{M}}(\beta)$ must be finite. Furthermore, $\mathbf{z}^{\mathrm{M}}(\beta) \succ \mathbf{z}^{\mathrm{M}}(\beta')$ for all β, β' such that $\underline{\beta} < \beta < \beta' < \widetilde{\beta}$.*

[13] The function $p \mapsto p\overline{D}(\zeta(p))$ is equivalent to the "pressed" $\overline{D}$ introduced by Libgober and Mu (2021).

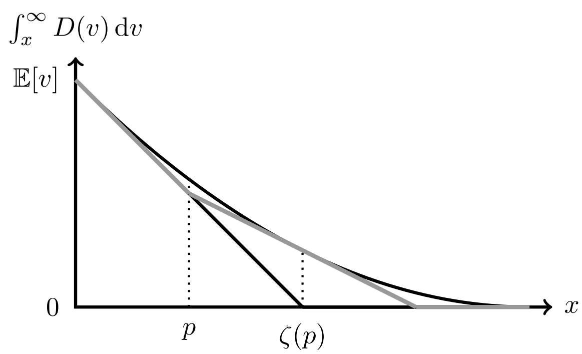

Figure 5: Minimizing Revenue Given p

The welfare implication of Proposition 2 is essentially the same as that of Proposition 1. The only differences are (i) the upper bound below which a unique Markov perfect equilibrium outcomes exists becomes lower when infinite Markov perfect equilibria are allowed, and (ii) the range of market feedback where unintended welfare loss is caused by higher market feedback becomes smaller.

Similar to Proposition 1, we prove Proposition 2 by characterizing every infinite Markov perfect equilibria. To state this characterization, for any $q \in [0,1]$, define

$$S(q) := \int_0^q \overline{D}^{-1}(z)\,\mathrm{d}z.$$

Note that $S(\overline{D}(p)) = p\overline{D}(p) + \int_p^\infty \overline{D}(v)\,\mathrm{d}v$ is the sum of consumer surplus and sales revenue when the price is p and when the demand is $\overline{D}$. With this additional notation, we can now state the characterization.

Theorem 2 (Markov Perfect Equilibrium Outcomes)**.** *For any $\beta \geq 0$, an infinite Markov perfect equilibrium exists if and only if $\beta \geq \widetilde{\beta}$. Moreover, for any Markov perfect equilibrium outcome $\mathbf{z}^{\mathrm{M}} = (r^{\mathrm{M}}, \sigma^{\mathrm{M}}, \omega^{\mathrm{M}}, p^{\mathrm{M}}, \{m_t^{\mathrm{M}}\})$, exactly one of the following is true.*

1. *$\mathbf{z}^{\mathrm{M}}$ is a finite Markov perfect equilibrium outcome.*

2. *$\omega^{\mathrm{M}} = \infty$,*

$$\begin{aligned} \frac{1-\gamma\delta}{\delta\beta} &\leq \sigma^{\mathrm{M}} \leq \mathbb{E}[v] - r^* \\ \mathbb{E}[v] - \sigma^{\mathrm{M}} &\leq p^{\mathrm{M}} \leq \frac{r^{\mathrm{M}}}{S^{-1}(r^{\mathrm{M}} + \sigma^{\mathrm{M}})}, \\ \max\left\{S(\overline{D}(\zeta(\mathbb{E}[v] - \sigma^{\mathrm{M}}))) - \sigma^{\mathrm{M}}, r^*\right\} &\leq r^{\mathrm{M}} \leq \mathbb{E}[v] - \sigma^{\mathrm{M}} \end{aligned}$$

and $m_t^{\mathrm{M}} = (\gamma + \beta\sigma^{\mathrm{M}})^t$ for all $t \geq 1$.

5 Equilibrium Payoff Set and the Efficiency Lower Bound

In this section, we move beyond Markov perfect equilibria and characterize the set of payoffs that can be supported by a subgame perfect equilibrium. Due to the infinitely repeated nature, there are inevitably many subgame perfect equilibrium outcomes. The results in this section characterize these outcomes, which in turn lead to another version of unintended welfare loss. Indeed, as shown below, whenever $\beta \leq \overline{\beta}$, there exists a unique *Pareto-worst* subgame perfect equilibrium outcome. Moreover, this worst outcome becomes Pareto inferior as β increases.

To state the results, first recall that r^* denotes the revenue guarantee defined in (8). Furthermore, since the function $p \mapsto p\overline{D}(v^{-1}(p))$ is quasi-concave, there exists a unique $p^* \geq \mathbb{E}[v]$ such that $r^* = p^*\overline{D}(v^{-1}(p^*))$. In other words, when the seller charges a price p^* and when the intermediary chooses her myopic best response given p^*, the sales revenue would be exactly r^*. Consequently, in any subgame perfect equilibrium, the seller would not charge any price $p > p^*$, since the best revenue given that price is below the revenue guarantee r^*.[14] With the definition of r^* and p^*, define

$$\omega^* := \frac{\alpha p^* \overline{D}(v^{-1}(p^*))}{1-\gamma\delta} = \frac{\alpha r^*}{1-\gamma\delta},$$

and notice that ω^* is the present discounted profit of the intermediary if the market growth rate is γ (i.e., consumer surplus is zero) and the sales revenue equals to the revenue guarantee in every period. Meanwhile, let $\hat{p}$ be the unique price $p \leq \overline{p}$ such that $p\overline{D}(p) = r^*$, define

$$\widehat{\beta} := \frac{1-\gamma\delta}{\delta\int_{\hat{p}}^{\infty}\overline{D}(v)\,\mathrm{d}v} \quad \text{and} \quad \beta^* := \frac{1-\gamma\delta}{\delta\int_{p^*}^{\infty}\overline{D}(v)\,\mathrm{d}v}$$

and notice that $0 < \underline{\beta} < \widetilde{\beta} < \widehat{\beta} < \overline{\beta} < \beta^*$.

In what follows, we will characterize the set of intermediary's payoffs among all subgame perfect equilibria. To begin with, we first identify a useful lower bound of the intermediary's equilibrium payoff. Clearly, since the sales revenue in every period is at least r^* and since the market growth rate is at least γ, the intermediary's payoff must be at least ω^*. Moreover, if the seller prices at p^* in every period and the intermediary chooses the myopic best response, then the payoff ω^* can be attained. However, it is not always incentive compatible for the seller and the intermediary to adopt these strategies. In particular, even when the intermediary anticipates that the seller will always post p^* and that the normalized continuation value will always be ω^*, it may not be optimal for the intermediary to choose the myopic best response and maximize sales revenue when the market feedback is large enough. To take this into account, we construct a tighter lower bound and accounts for the aforementioned incentive.

[14]In fact, p^* plays the role as the *min-max strategy* does in canonical repeated games with perfect monitoring, and coincides with the min-max strategy in Fudenberg, Kreps, and Maskin (1990).

Let

$$h(\omega) := \delta\left(\gamma + \beta\int_{\left(1-\frac{\alpha}{\delta\beta\omega}\right)^{+}p^*}^{\infty} \overline{D}(v)\,\mathrm{d}v\right)\omega,$$

for all $\omega \geq 0$, and let

$$\underline{\omega} := \inf\{\omega \geq \omega^* | h(\omega') \leq 1,\, \forall \omega' \geq \omega\}.$$

Notice that $\underline{\omega} = \omega^*$ whenever $\beta \leq \overline{\beta}$, and that $\underline{\omega} \to \infty$ as $\beta \uparrow \beta^*$.

Lemma 5. *In any subgame perfect equilibrium, the intermediary's payoff is at least* $\underline{\omega}$.

Similar to Section 4, we first characterize the finite subgame perfect equilibrium outcomes (i.e., subgame perfect equilibrium outcomes where the intermediary's continuation value is finite in every subgame). For any $\beta \geq 0$, let $\Omega^f(\beta)$ denote the set of intermediary's payoffs among all finite subgame perfect equilibrium outcomes.

Theorem 3 (Finite Subgame Perfect Equilibrium Payoffs)**.** *A finite subgame perfect equilibrium exists if and only if* $\beta < \beta^*$*. Furthermore, for any* $\beta \in [0, \beta^*)$*,*

$$\Omega^f(\beta) = [\underline{\boldsymbol{\omega}}^f(\beta), \overline{\boldsymbol{\omega}}^f(\beta)] \backslash \{\infty\},$$

for some $\underline{\omega} \leq \underline{\boldsymbol{\omega}}^f(\beta) \leq \overline{\boldsymbol{\omega}}^f(\beta) \leq \infty$*. Moreover,* $\underline{\boldsymbol{\omega}}^f$ *is nonincreasing on* $[0, \overline{\beta}]$ *and* $\overline{\boldsymbol{\omega}}^f$ *is nondecreasing on* $[0, \beta^*)$*; while* $\underline{\boldsymbol{\omega}}^f(\beta) = \underline{\omega}$ *whenever* $\beta \in [\widehat{\beta}, \beta^*)$*; and*

$$\overline{\boldsymbol{\omega}}^f(\beta) = \begin{cases} \frac{\alpha\mathbb{E}[v]}{1-\gamma\delta}, & \textit{if } \beta \in [0, \underline{\beta}] \\ \infty, & \textit{if } \beta \in [\widehat{\beta}, \beta^*) \end{cases}.$$

As a remark, Theorem 3 characterizes the long run player's equilibrium payoff for every fixed discount $\delta \in (0, 1)$, rather than only characterizing these payoffs as δ approaches to one as in Fudenberg, Kreps, and Maskin (1990). Furthermore, the characterization does not involve fixed-point arguments and the notion of self-generating as in Abreu, Pearce, and Stacchetti (1986) and Abreu, Pearce, and Stacchetti (1990).[15] Instead, the bounds $\underline{\boldsymbol{\omega}}^f(\beta)$ and $\overline{\boldsymbol{\omega}}^f(\beta)$ are simply determined by the value of a constrained optimization problem.

In addition to characterizing the intermediary's equilibrium payoff, it would also be useful to explore the sales revenue, consumer surplus, and prices that may occur in any period under a finite subgame perfect equilibrium. After all, knowing how consumer surplus and sales revenue are structured in a subgame perfect equilibrium could provide further welfare implications. To this end, for any $\beta \geq 0$, let

$$\underline{r}(\beta) := \begin{cases} \frac{(1-\gamma\delta)\underline{\boldsymbol{\omega}}^f(\beta)}{\alpha}, & \text{if } \beta \leq \widehat{\beta}. \\ r^*, & \text{if } \beta > \widehat{\beta} \end{cases}$$

[15]The reason is that—in addition to perfect monitoring—the intermediary and the seller's stage game payoffs are linearly dependent, so that feasible payoffs in a stage game is a line segment. As a result, characterizing the intermediary's equilibrium payoff is essentially equivalent to finding two endpoints.

and define the set $\mathbf{Z}^f(\beta)$ as follows:

$$\mathbf{Z}^f(\beta) := \left\{ (r,\sigma,p) \in \mathbb{R}^3_+ \;\middle|\; \begin{array}{c} (\mathbb{E}[v]-p)^+ \leq \sigma \leq S\left(\frac{r}{p}\right) - r \\ \underline{r}(\beta) \leq r \leq p \\ \alpha r + \delta(\gamma+\beta\sigma)\overline{\boldsymbol{\omega}}^f(\beta) \geq \sup_{D\in\mathcal{D}} \left[\alpha p D(p) + \delta\left(\gamma + \beta \int_p^\infty D(v)\,\mathrm{d}v\right)\underline{\boldsymbol{\omega}}^f(\beta)\right] \end{array} \right\}$$

Using Theorem 3, finite subgame perfect equilibrium outcomes can be characterized, and is given by Corollary 1

Corollary 1 (Finite Subgame Perfect Equilibrium Outcomes)**.** *For any $\beta \geq 0$ and for any finite subgame perfect equilibrium outcome $\mathbf{z} = \{r_t, \sigma_t, \omega_t, p_t, m_t\}$, $(r_t, \sigma_t, p_t) \in \mathbf{Z}^f(\beta)$ for all $t \in \mathbb{N}$. Furthermore, for any $\beta \geq 0$, for any $T \in \mathbb{N}$, and for any $(r,\sigma,p) \in \mathbf{Z}^f(\beta)$, there exists a finite subgame perfect equilibrium outcome $\mathbf{z} = \{r_t, \sigma_t, \omega_t, p_t, m_t\}$ such that $r_T = r$, $\sigma_T = \sigma$, and $p_T = p$.*

An immediate consequence of Corollary 1 is that any outcome $\mathbf{z} = \{r_t, \sigma_t, \omega_t, p_t, m_t\}$ in which $r_t = \underline{r}(\beta)$, $\sigma_t = 0$, and $\omega_t = \underline{\boldsymbol{\omega}}^f(\beta)$ for all t is dominated by any other equilibrium outcomes. In fact, whenever $\beta \leq \overline{\beta}$, such an equilibrium outcome always exists. In this equilibrium, the seller always charges a price $\underline{p} \in [\mathbb{E}[v], p^*]$ so that $\underline{p}\overline{D}^{-1}(v^{-1}(\underline{p})) = \underline{r}(\beta)$, and the intermediary always chooses her myopic best response, which leads to zero consumer surplus in each period. In the event of any deviation from the seller, the intermediary chooses a disclosure policy to minimize revenue and the tie-breaker breaker breaks tie against the seller. If the intermediary follows this punishment, then the continuation play gives an equilibrium payoff of $\overline{\boldsymbol{\omega}}^f(\beta)$; otherwise the continuation play gives the intermediary payoff $\underline{\boldsymbol{\omega}}^f(\beta)$.[16]

As a result, whenever $\beta \leq \overline{\beta}$, there exists a subgame perfect equilibrium outcome that is dominated by every other subgame perfect equilibrium outcomes. Moreover, since $\underline{\boldsymbol{\omega}}$ is nonincreasing in β, this dominated outcome becomes Pareto-inferior as β increases.

Proposition 3 (Unintended Welfare Loss—Finite Subgame Perfect)**.** *For any $\beta \in [0, \overline{\beta}]$, there exists a finite subgame perfect equilibrium outcome $\mathbf{z}^*(\beta)$ that is dominated by any other finite subgame perfect equilibrium outcomes. Furthermore, for any γ, δ such that $\gamma\delta \leq 1/2$, there exists $\widehat{\beta}(\gamma,\delta) \in (0, \widehat{\beta})$ such that for any $0 < \beta < \beta' < \widehat{\beta}(\gamma,\delta)$, $\mathbf{z}^*(\beta) \succ \mathbf{z}^*(\beta')$.*

It is noteworthy that the worst subgame perfect equilibrium outcome $\mathbf{z}^*(\beta)$ introduced by Proposition 3 is not Markov. In essence, $\mathbf{z}^*(\beta)$ is obtained by constructing an equilibrium that incentivizes the seller to post a high price, under which the intermediary's best response, when anticipating the seller posting the same high price in the future, is to maximize the

[16]By the definition of $\underline{\boldsymbol{\omega}}^f(\beta)$ and $\overline{\boldsymbol{\omega}}^f(\beta)$ (which can be found in the online appendix), this makes the intermediary's punishment incentive compatible.

current sales revenue and leave zero surplus to consumers. To properly incentivize the seller, the intermediary punishes the seller's deviation by minimizing sales revenue. For such punishments to be incentive compatible for the intermediary, the reward for carrying out the punishment has to be large enough. Higher market feedback β allows the market to grow faster and hence allows for higher continuation payoff for rewarding. As a result, when β is higher, more severe punishments can be supported and hence more extreme price can be incentivized.

We conclude this section by discussing subgame perfect equilibria where the intermediary's payoff may be infinite in some subgames. From Theorem 2, it follows immediately that an infinite subgame perfect equilibrium exists whenever $\beta \geq \widetilde{\beta}$. In fact, the converse is also true: An infinite subgame perfect equilibrium exists only if $\beta \geq \widetilde{\beta}$. When an infinite subgame perfect equilibrium exists, the intermediary can always be incentivized to use the most severe punishment in the event of a deviation, and hence the seller can be incentivized to charge any price in $[r^*, p^*]$. As a result, any feasible payoff can be supported by a subgame perfect equilibrium, as summarized below. To this end, for any $\beta \geq 0$, let $\Omega(\beta)$ denote the set of intermediary's payoffs among all subgame perfect equilibria.

Theorem 4 (Subgame Perfect Equilibrium Payoffs)**.** *An infinite subgame perfect equilibrium exists if and only if $\beta \geq \widetilde{\beta}$. Furthermore, for any $\beta \geq 0$,*

$$\Omega(\beta) = \begin{cases} \Omega^f(\beta), & \text{if } \beta < \widetilde{\beta} \\ [\underline{\omega}, \infty], & \text{if } \beta \geq \widetilde{\beta} \end{cases}.$$

In the meantime, for any $\beta \geq 0$, let $\mathbf{Z}(\beta)$ be defined as

$$\mathbf{Z}(\beta) := \begin{cases} \mathbf{Z}^f(\beta), & \text{if } \beta < \widetilde{\beta} \\ \left\{ (r, \sigma, p) \in \mathbb{R}^3_+ \,\middle|\, \begin{array}{c} r^* \leq r \leq p \\ (\mathbb{E}[v] - p)^+ \leq \sigma \leq S\left(\frac{r}{p}\right) - r \end{array} \right\}, & \text{if } \beta \geq \widetilde{\beta} \end{cases}.$$

Then the subgame perfect equilibrium outcomes can be characterized as well.

Corollary 2 (Subgame Perfect Equilibrium Outcomes)**.** *For any $\beta \geq 0$ and for any subgame perfect equilibrium outcome $\mathbf{z} = \{r_t, \sigma_t, \omega_t, p_t, m_t\}$, $(r_t, \sigma_t, p_t) \in \mathbf{Z}(\beta)$ for all $t \geq 0$. Furthermore, for any $\beta \geq 0$, for any $T \geq 0$, and for any $(r, \sigma, p) \in \mathbf{Z}(\beta)$, there exists a subgame perfect equilibrium outcome $\mathbf{z} = \{r_t, \sigma_t, \omega_t, p_t, m_t\}$ such that $r_T = r$, $\sigma_T = \sigma$, and $p_T = p$.*

Together with Corollary 2, the implication of Proposition 3 can be extended even when allowing for infinite subgame perfect equilibria, as summarized below.

Proposition 4 (Unintended Welfare Loss—Subgame Perfect)**.** *For any $\beta \in [0, \overline{\beta}]$, there exists a subgame perfect equilibrium outcome $\mathbf{z}^*(\beta)$ that is dominated by any other subgame perfect equilibrium outcomes. Furthermore, for any γ, δ such that $\gamma\delta \leq 1/2$, there exists $\widetilde{\beta}(\gamma, \delta) \in (0, \widetilde{\beta}]$ such that for any $0 < \beta < \beta' < \widetilde{\beta}(\gamma, \delta)$, $\mathbf{z}^*(\beta) \succ \mathbf{z}^*(\beta')$.*

6 Policy Implications: Countering the Unintended Welfare Loss

As demonstrated by Proposition 1 (as well as Proposition 3), higher level of market feedback may lead to Pareto inferior outcomes. The underlying driving force of this result is the difference in interest between the seller and the intermediary. More specifically, Proposition 1 is driven by the fact that the intermediary would prefer the consumers to have larger surplus *and* the seller to have higher revenue at the same time. As a result, as the level of market feedback increases, the seller can raise the price to compel the intermediary to provide information in a way that is more aligned with the seller's interest. Based on this observation, we now discuss several alternatives that could alleviate the problem of unintended welfare loss: subscription-based model, capping the market growth, and long-lived seller.

One of these alternatives involves a different business model for the intermediary. Specifically, instead of sharing the sales revenue with the seller, the intermediary's source of revenue can be derived from the consumers. Under the *subscription-based* business model, by either collecting subscription fees from consumers directly or monetizing one's digital account, the intermediary's profit becomes independent of the seller's sales revenue—and in particular, seller's posted price. This in turn implies that the seller cannot affect the intermediary's disclosure policy using his posted price. In fact, as shown below, outcomes are always *more* efficient when the market feedback level is higher.

Another alternative pertains to controlling the level of market feedback by imposing a cap on the market growth rate.[17] By capping the market growth rate, the difference of interest between the intermediary and the seller can be further controlled. As shown below, under proper caps, higher market feedback level would not lead to Pareto-worse outcomes and would generate strictly higher total surplus in the economy than what is achieved without a cap.

The third alternative seeks to eliminate the difference in interest by establishing a long-term relationship between the seller and the intermediary. After all, the difference in interest is created by the feature that the seller does not internalize the benefit of market growth at the time of setting the price. By establishing a long-term relationship with the intermediary (i.e., by considering a long-lived seller), as shown below, higher market feedback would never lead to less efficient outcomes.

6.1 Subscription-Based Model

In this section, we consider an alternative business model. In each period $t \geq 0$, the timing and the strategies remain the same: The seller posts a price p_t, the intermediary sees p_t and

[17] In digital markets, this can be achieved by properly designing the recommendation algorithms so that an intermediary cannot be targeted to too many consumers in a short period of time.

chooses a disclosure policy $D_t \in \mathcal{D}$, and the tie-breaker sees p_t and D_t and chooses a tie-breaking rule $q_t \in [D_t(p_t^+), D_t(p_t)]$. However, instead of sharing with the intermediary, the seller captures all the sale revenue $m_t \cdot p_t \cdot q_t$ in each period. The intermediary, on the other hand, receives $\widetilde{\alpha} \in [0,1]$ share of the consumer surplus in each period. The interpretation is that upon arrival, consumers pay a share $\widetilde{\alpha}$ of their ex-ante surplus to the intermediary in exchange of the intermediary's information.

In essence, the subscription-based model decouples the seller's and the intermediary's interests. Rather than seeking to raise sales revenue while keeping consumers surplus high enough to sustain future revenue, the intermediary's only goal in the subscription-based model is to maximize consumer surplus, as both her stage game payoff and the market size depend only on consumers' average surplus. This further leads to a way to circumvent the unintended welfare loss.

Proposition 5 (Subscription-Based Model)**.** *A Markov perfect equilibrium always exists. For any $\beta < \overline{\beta}$, there exists a unique finite Markov perfect equilibrium outcome $\mathbf{y}^{\mathrm{M}}(\beta)$. Furthermore, for any $0 < \beta < \beta'$, $\mathbf{y}^{\mathrm{M}}(\beta) \prec \mathbf{y}^{\mathrm{M}}(\beta')$.*

Although it is clear from Proposition 1 and Proposition 5 that compared to the revenue-sharing model, the subscription-based model can better translate higher level of market feedback into more efficient outcomes, it is not clear that the intermediary would always prefer the subscription based model. After all, while the consumers enjoy higher surplus, and hence market sizes are large under subscription-based model, more surplus is extracted from the consumers under the revenue-sharing model. As demonstrated by Proposition 6 below, it is possible that the intermediary would always prefer the revenue-sharing model, even if it means that there would be unintended welfare loss.

To state this result, for any $\beta < \overline{\beta}$, let $\omega^{\mathrm{M}}(\beta)$ denote the intermediary's payoff in the unique finite Markov perfect equilibrium under the revenue-sharing model; and let $\rho^{\mathrm{M}}(\beta)$ denote the intermediary equilibrium in the unique finite Markov perfect equilibrium under the subscription-based model.

Proposition 6. *There exists $\beta^0 \geq 0$ such that $\omega^{\mathrm{M}}(\beta) > \rho^{\mathrm{M}}(\beta)$ for all $\beta \in [0, \beta^0)$. Moreover, $\beta^0 > 0$ if and only if*

$$\frac{\widetilde{\alpha}}{\alpha} < \frac{\mathbb{E}[v]}{\int_{\overline{p}}^{\infty} \overline{D}(v)\,\mathrm{d}v};$$

and $\beta^0 > \underline{\beta}$ if and only if

$$\frac{\widetilde{\alpha}}{\alpha} + 1 < \frac{\mathbb{E}[v]}{\int_{\overline{p}}^{\infty} \overline{D}(v)\,\mathrm{d}v}.$$

From Proposition 6, an immediate implication is the possibility that the intermediary would prefer the revenue-sharing model over the subscription-based model, even if the β falls in the range ($\beta > \underline{\beta}$) where higher market feedback level leads to unintended welfare loss under the revenue-sharing model. Consequently, while the subscription-based model may be better in terms of translating higher market feedback into efficiency, the intermediary might not be willing to adopt this model voluntarily. Hence, some third-party interventions (e.g., monetization by large platforms and prohibiting revenue sharing) would be necessary.

6.2 Capping the Market Growth

Besides interventions that aim to alter the intermediary's business model, we can also avoid the unintended welfare loss through policy instruments that affect the market feedback. In this section, we consider an alternative setting in which regulators can impose a cap on the market growth rate. In each period, the timing, the strategies, and the payoffs remain the same. However, the growth rate of the market size in each period cannot exceed an upper bound χ. As a result, the evolution of market sizes in period $t \geq 0$ is

$$m_{t+1} = m_t \cdot \min\left\{\chi, \gamma + \beta \int_{p_t}^{\infty} D_t(v)\,\mathrm{d}v\right\}.$$

Intuitively, with a cap on the market growth rate, market feedback is effectively zero once the market growth rate reaches the cap χ and the intermediary therefore has no incentive to further raise the consumer surplus. Hence, a cap on the market growth rate limits the extent to which the intermediary cares about consumer surplus and aligns the intermediary's interests with the seller's. As the cap on market growth rate gets tighter, the intermediary's interests get better aligned with the seller's. This leads to another way to circumvent the unintended welfare loss.

Proposition 7. *For any $\beta \geq 0$, there exist $\overline{\chi} > \gamma$ such that for any $\chi \in (\gamma, \overline{\chi})$, there is a unique Markov perfect equilibrium outcome $\mathbf{x}^{\mathrm{M}}(\beta; \chi)$ and $m_t = \chi^t$ in equilibrium.*

For each $\beta \geq 0$, let $\overline{\chi}^{\beta}$ denote the supremum of the set of $\overline{\chi}$ for which the statement in Proposition 7 is valid. For $\chi \in (\gamma, \overline{\chi}^{\beta})$, let $S_t^{\mathrm{M}}(\beta; \chi)$ denote the total surplus (the sum of the seller's, the intermediary's, and the consumers' payoffs) in period t under $\mathbf{x}^{\mathrm{M}}(\beta; \chi)$. Let $S_t^{\mathrm{M}}(\beta; \infty)$ denote the total surplus in period t under the equilibrium outcome $\mathbf{z}^{\mathrm{M}}(\beta)$ in the baseline model without a cap on market growth.

Proposition 8. *(Capping the Market Growth) For any $\underline{\beta} < \beta_1 < \beta_2 < \overline{\beta}$ and any $\chi \in (\gamma, \min\{\overline{\chi}^{\beta_1}, \overline{\chi}^{\beta_2}\})$,*

1. $\mathbf{x}^{\mathrm{M}}(\beta_1; \chi) \not\succ \mathbf{x}^{\mathrm{M}}(\beta_2; \chi)$.
2. $S_t^{\mathrm{M}}(\beta_1; \chi) = S_t^{\mathrm{M}}(\beta_2; \chi) = \chi^t \mathbb{E}(v) > S_t^{\mathrm{M}}(\beta_1; \infty) > S_t^{\mathrm{M}}(\beta_2; \infty)$.

As demonstrated in Proposition 8, a cap on the market growth rate can ensure that higher market feedback does not lead to an unintended Pareto-worse outcome and can generate a total surplus in each period strictly higher than the total surplus when there is no cap on market growth. Intuitively, the intermediary has lower commission fee per unit of sales and cares more about generating market growth when the price is lower. Therefore, the cap on the growth rate only binds at lower prices, in which case the intermediary's interests get aligned to the seller's. Anticipating this, the seller lowers the price and induces more consumers to purchase the product. In equilibrium, all gains from trade are realised and the market size grows faster, both of which contributes to a higher total surplus.

Besides raising the total surplus, a cap on the market growth rate can also lead to Pareto improvements *in the long run* if there is non-negative exogenous growth ($\gamma > 1$). To state the result, let $\mathbf{x}_t^{\mathrm{M}}(\beta, \chi)$ and $\mathbf{z}_t^{\mathrm{M}}(\beta)$ denote the *period-t outcomes* under $\mathbf{x}^{\mathrm{M}}(\beta, \chi)$ and $\mathbf{z}^{\mathrm{M}}(\beta)$ respectively. Period-t outcomes can also be Pareto-ranked according to Definition 1.

Proposition 9. *Suppose that $\gamma > 1$. For any $\beta \in (\underline{\beta}, \overline{\beta})$, there exists $\overline{\chi} > \gamma$ and $\bar{t} > 0$ such that $\mathbf{x}_t^{\mathrm{M}}(\beta; \chi) \succ \mathbf{z}_t^{\mathrm{M}}(\beta)$ for $\chi \in (\gamma, \overline{\chi})$ and $t > \bar{t}$.*

Even though a cap on the market growth rate that is tight enough can circumvent the unintended welfare loss, generate additional surplus, and make all market participants better off in the long run when $\gamma > 1$, restricting market growth too much would lead to a slow market growth and reduced surplus in the market, as shown in Proposition 8.

Proposition 10. *For any $\beta > \underline{\beta}$, $S_t^{\mathrm{M}}(\beta; \cdot)$ is strictly increasing on $(\gamma, \overline{\chi}^{\beta})$.*

Consequently, policy makers should carefully choose the imposed cap on the market growth rate to ensure that it is below the threshold $\overline{\chi}^{\beta}$ but not too much lower.

6.3 Long-Lived Seller

We now consider an alternate model where the seller is long-lived. The seller being long-lived can be interpreted as establishing a long-term relationship between the seller and the intermediary, which allows the seller to share the benefit of market growth and thus be market-minded as well. Specifically, suppose that the seller is long-lived and that both the seller and the intermediary have discount $\delta \in (0, 1)$. The timing and strategies of all players remain the same: In each period t, the seller observes all past histories and then chooses a price p_t; the intermediary then sees all past histories and p_t before choosing a disclosure policy $D_t \in \mathcal{D}$; the tie-breaker then sees all past histories, p_t and D_t and chooses a tie breaking rule $q_t \in [D_t(p_t^+), D_t(p_t)]$. Given any strategy profile, the seller's payoff is

$$\pi = (1-\alpha)p_0 q_0 + \sum_{t=1}^{\infty} \delta^t \prod_{s=0}^{t} \left(\gamma + \beta \int_{p_s}^{\infty} D_s(v)\, \mathrm{d}v \right) (1-\alpha) p_t q_t,$$

while the intermediary's payoff is

$$\omega = \alpha p_0 q_0 + \sum_{t=1}^{\infty} \delta^t \prod_{s=0}^{t} \left(\gamma + \beta \int_{p_s}^{\infty} D_s(v)\,\mathrm{d}v \right) \alpha p_t q_t,$$

where $\{p_t, D_t, q_t\}$ are the actions chosen by the seller, the intermediary, and the tie-breaker on path, respectively. As a result, the seller's and the intermediary's interests are perfectly aligned in this alternate model. As demonstrated by Proposition 11 below, this prevents unintended welfare losses as β increases.

Proposition 11 (Long-Lived Seller)**.** *A Markov perfect equilibrium always exists. A finite Markov perfect equilibrium outcome exists if and only if $\beta < \underline{\beta}$. Moreover, whenever $\beta < \underline{\beta}$, the Markov perfect equilibrium outcome is unique and is the same as that when the seller is short-lived.*

From Proposition 11, it then follows that when $\beta \leq \underline{\beta}$, then whether the seller and the intermediary have a long-run relationship is irrelevant. The induced Markov perfect equilibrium outcome is the same. In contrast, when $\beta > \underline{\beta}$, from Proposition 1 and Proposition 2, unintended welfare loss is possible when the seller is short-lived. However, the finite Markov perfect equilibria that lead to unintended welfare losses would not even exist if the seller is long-lived. This suggests that a possible way to avoid unintended welfare losses is to encourage sellers and the intermediaries to form long-term relationships and to discourage sellers from frequently switching intermediaries they collaborate with.

7 Extensions

7.1 Strategic Consumers

Throughout the paper, we have abstracted away consumers participation decisions and assumed an exogenous law of motion of the market size. In reality, consumers could decide whether they would like to participate and become a potential customer of the intermediary. A certain aspect of this feature can be incorporated into the baseline model, which we now describe.

Consider the baseline model but suppose instead that the market size evolves as follows: In every period $t \geq 1$, a mass $\gamma^c m_{t-1}$ of *captive* consumers, and a mass $\beta^c m_{t-1}$ of *shopping* consumers first arrive, for some $\gamma^c, \beta^c \geq 0$. The captive consumers must remain as a potential customer of the intermediary (and hence the seller) regardless, while the shopping consumers have outside options that are uniformly distributed on $[0, \mathbb{E}[v]]$, and can decide whether to leave the market. If s_t is the share of shopping consumers who decide not to leave, then the market size in period t would be

$$m_t = m_{t-1}(\gamma^c + \beta^c s_t).$$

In this alternate setting, it follows that if the seller, the intermediary, and the tie-breaker play any Markov strategy profile in the baseline model with $\gamma = \gamma^c$ and $\beta = \beta^c/\mathbb{E}[v]$, then a best response for shopping consumers is to leave the market if and only if their outside option is greater than their expected surplus induced by the strategy profile. Meanwhile, when consumers follow this strategy, the market size evolution becomes

$$m_t = m_{t-1}\left(\gamma^c + \frac{\beta^c}{\mathbb{E}[v]}\int_{p_t}^{\infty} D_t(v)\,\mathrm{d}v\right)$$

for all t, where $p_t \geq 0$ and $D_t \in \mathcal{D}$ are the price and the disclosure policy chosen by the seller and the intermediary on path under that strategy profile, respectively.

As a result, any Markov perfect equilibrium in the baseline model with $\gamma = \gamma^c$ and $\beta = \beta^c/\mathbb{E}[v]$ would correspond to an equilibrium in this alternate setting where some consumers can choose to leave.[18] Moreover, consumers' strategies in this equilibrium is simple: they leave if and only if their outside option is greater than the average surplus in the previous period.

7.2 Nonlinear Growth Rate

In the baseline model, it is assumed that the market growth rate is linear in consumer surplus. In practice, market growth rates could depend on consumer surplus in a non-linear manner. In this section, we consider an extension that captures a broad class of non-linear market growth. Specifically, in each period $t \geq 1$, the growth rate of the market size m_{t+1}/m_t can be a nonlinear function of the consumer surplus in period t:

$$\frac{m_{t+1}}{m_t} = f\left(\int_{p_t}^{\infty} D_t(v)\,\mathrm{d}v\right).$$

for some function f, which we refer to as the market growth function. In this section, assume that the function $p \mapsto p\overline{D}(p)$ is strictly concave. We now characterize the set of Markov perfect equilibria in this alternative setting.

Let $\mathcal{F}$ be the collection of twice-differentiable, increasing, and concave functions on $\mathbb{R}_+$ with $f(0) = \gamma$, and let

$$\mathcal{F}_1 := \{f \in \mathcal{F} | f'(0) \in [0, \underline{\beta}]\}.$$

Furthermore, for any $\beta \in (\underline{\beta}, \overline{\beta})$ and for any $\eta \geq 0$, let

$$\mathcal{F}_2(\beta, \eta) := \{f \in \mathcal{F} | f'(0) = \beta,\ \|f''\| \leq \eta\}.$$

[18]This also provides a foundation for the interpretation that market feedback is higher when competition is more intense, as the intensity of competition among intermediaries can be captured by the parameter β^c.

Notice that for any $\beta \in (\underline{\beta}, \overline{\beta})$, $\mathcal{F}_1$ and $\mathcal{F}_2(\beta, \eta)$ are disjoint. Moreover, in comparison to the baseline model where market growth rates are assumed to be an affine function of consumer surplus, $f \in \bigcup_{\beta \in (\underline{\beta}, \overline{\beta})} \mathcal{F}_2(\beta, 0)$ is equivalent to the case of $\beta \in (\underline{\beta}, \overline{\beta})$.

As shown in Proposition 12 below, even with nonlinear market growth rate, there may be unintended welfare losses as the level of market feedback increases.

Proposition 12 (Unintended Welfare Loss—Nonlinear Growth)**.** *There exists a continuously decreasing function $h : (\underline{\beta}, \overline{\beta}) \to \mathbb{R}_{++}$ such that every $f \in \mathcal{F}_1 \cup [\bigcup_{\beta \in (\underline{\beta}, \overline{\beta})} \mathcal{F}_2(\beta, h(\beta))]$ induces a unique finite Markov perfect equilibrium outcome $\mathbf{z}^{\mathrm{M}}(f)$. Furthermore, for any β, β' such that $\underline{\beta} < \beta < \beta' < \overline{\beta}$, $\mathbf{z}^{\mathrm{M}}(f_1) \succ \mathbf{z}^{\mathrm{M}}(f_2)$ for all $f_1 \in \mathcal{F}_2(\beta, h(\beta))$ and $f_2 \in \mathcal{F}_2(\beta', h(\beta'))$.*

According to Proposition 12, for any growth function in the set $\bigcup_{\beta \in (\underline{\beta}, \overline{\beta})} \mathcal{F}_2(\beta, h(\beta))$, an increase in the level of market feedback at the level of zero consumer surplus (i.e., $f'(0)$) would unambiguously lead to welfare loss. From this regard, the unintended welfare losses are not exclusively implied by linearity of the market growth function. Rather, the linearity assumption in the baseline model is largely for the ease of exposition. Moreover, according to Proposition 12, except for the derivative at zero $f'(0)$ and the second derivative $\|f''\|$, behaviors of growths functions $f \in \mathcal{F}$ do not matter for our result. Instead, only the second derivative $\|f''\|$ needs to be disciplined as $f'(0)$ increases (through the function h).

7.3 Non-Stationary Revenue-Sharing Rule

Thus far, it is assumed that the seller and the intermediary share sales revenues in each period according to a fixed $\alpha \in (0, 1)$. It is, however, reasonable to expect non-constant sharing rules in reality. After all, in a market consists of many intermediaries and many sellers, intermediaries with different prominence levels may have different outside options, and are in fact offering different products (since the mass of consumers that can be reached would be different). In this section, we consider an extension where the sharing rule can depend on the current market size m, so that the intermediary can obtain $\boldsymbol{\alpha}(m)$ share of the sales revenue when the market size is m. When the sharing rule depends on the market size, the stage game is no longer stationary in general. As a result, Markov strategies may depend on market sizes in general and hence the intermediary's best response in each period would be characterized by a recursive Bellman equation even when restricted attention to finite Markov perfect equilibria. Nonetheless, as shown below, under certain parametrization of the function $\boldsymbol{\alpha}$, our previous analyses can be readily extended.

In what follows, we assume that at any history where the market size is $m \geq 1$, the intermediary can retain $\boldsymbol{\alpha}(m) := m^{-\alpha}$ share of the sales revenue, for some $\alpha \in (0, 1)$.[19] With

[19]Under this assumption, $\boldsymbol{\alpha}$ is decreasing in m, which means that the intermediary's per-consumer per-sale commission is decreasing in its prominence. Although under a different model, the equilibrium prices for

this assumption, under any strategy profile, for any $t \geq 0$, notice that

$$\widetilde{m}_{t+1} := m_{t+1}\boldsymbol{\alpha}(m_{t+1}) = f(\sigma_t)m_t\boldsymbol{\alpha}(m_t) =: f(\sigma_t)\widetilde{m}_t,$$

where $f(\sigma) := (\gamma + \beta\sigma)^{\alpha+1}$ for some $\gamma \geq 1$ is and σ_t denotes the consumer surplus induced by this strategy profile in period t. Note that f is an increasing and concave function and hence we may simply replace $\{m_t\}$ by $\{\widetilde{m}_t\}$ and apply the results in Section 7.2.

Thus, even if the revenue sharing rule between the seller and the intermediary is allowed to be non-stationary (in particular, to depend on the current market size), for a certain range of β there would still be unintended welfare loss as β increases.

8 Conclusions

In this paper, we show the possibility of unintended welfare losses in a setting where a product seller collaborates with an informational intermediary in a relatively short term and rewards intermediary via revenue sharing. The underlying reason is the difference in interests arising from the intermediary's concern about her own prominence in the future, and hence is more market-minded compared to the seller. Specifically, we show that for a range of market feedback levels, under the unique Markov perfect equilibria, higher market feedback would lead to Pareto worse outcomes—even if higher market feedback means that market would grow faster and that the intermediary is more market-minded and has greater incentives to enhance consumer surplus. Moreover, we show that these unintedned welfare losses may still be present even if one looks across all subgame perfect equilibria, in the sense that the Pareto-worst subgame perfect equilibrium becomes even worse as the market feedback increases.

We then discuss the implications of our characterization and the aforementioned unintended welfare losses. Our results can serve as a cautionary tale for policy making, as they highlight the possibility changes in the level of market feedback may have counter-intuitive policy implications. Moreover, we propose several ways that could eliminate such unintended welfare losses, including replacing the revenue-sharing business model with a subscription-based model; imposing a cap on the market growth rate; and establishing a long-term relationship between the seller and the intermediary.

Several directions appear naturally as future research questions. For example, while we only focus on the revenue-sharing and information provision business model, informational intermediaries often operate with multiple different models, including the pay-sponsorship model. It would be valuable to understand how intermediaries and the market choose among these models and what are the welfare implications. Alternatively, while our model only

intermediaries' services exhibit a similar feature in Fainmesser and Galeotti (2021).

considers horizontal information (e.g., product characteristics), vertical information (e.g., product quality) is sometimes also an crucial dimension of an intermediary's services. Lastly, from a design point of view, it would also be useful to better understand the broader impact of rating systems and consumers' communication network.

References

Abreu, D., D. Pearce, and E. Stacchetti (1986): "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, 39, 251–269.

——— (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.

Ali, S. N., N. Haghpanah, X. Lin, and R. Siegel (2022): "How to Sell Hard Information," *Quarterly Journal of Economics*, 137, 619–678.

Anderson, S. P. and R. Renault (2006): "Advertising Content," *American Economic Review*, 96, 93–113.

Armstrong, M. and J. Zhou (forthcoming): "Consumer Information and the Limits to Competition," *American Economic Review*.

Au, P. H. and M. Whitmeyer (2021): "Attraction Versus Persuasion," Working Paper.

Bar-Isaac, H. and J. Deb (2021): "Reputation with Opportunities for Coasting," *Journal of European Economic Association*, 19, 200–236.

Bergemann, D., B. Brooks, and S. Morris (2015): "The Limits of Price Discrimination," *American Economic Review*, 105, 921–957.

Bergemann, D., B. Brooks, and S. E. Morris (2017): "First Price Auctions with General Information Structures: Information for Bidding and Revenue," *Econometrica*, 85, 107–143.

——— (2021): "Search, Information, and Prices," *Journal of Political Economy*, 129, 2275–2319.

Bergemann, D. and M. Pesendorfer (2007): "Information Structures in Optimal Auctions," *Journal of Economic Theory*, 137, 580–609.

Biglaiser, G. (1993): "Middleman as Experts," *RAND Journal of Economics*, 24, 212–223.

Board, S. and J. Lu (2018): "Competitive Information Disclosure in Search Markets," *Journal of Political Economics*, 126, 1965–2010.

Boleslavsky, R., I. Hwang, and K. Kim (2019): "Competitive Advertising and Pricing," Working Paper.

Brooks, B. and S. Du (2021): "Optimal Auction Design with Common Values: An Informationally-Robust Approach," *Econometrica*, 89, 1313–1360.

Chakraborty, I., J. Deb, and A. Öry (2021): "When do Consumers Talk," Working Paper.

Che, Y.-K. and J. Hörner (2018): "Recommender System as Mechanisms for Social Learning," *Quarterly Journal of Economics*, 133, 871–925.

Chen, Y.-C. and X. Yang (2020): "Information Design in Optimal Auctions," Working Paper.

Deb, R. and A.-K. Roesler (2021): "Multi-Dimensional Screening: Buyer-Optimal Learning and Informational Robustness," Working Paper.

Doval, L. and V. Skreta (2021): "Purchase History and Product Personalization," Working Paper.

Du, S. (2018): "Robust Mechanisms Under Common Valuation," *Econometrica*, 86, 1565–1588.

Dworczak, P. and G. Martini (2019): "The Simple Economics of Optimal Persuasion," *Journal of Political Economy*, 127, 1993–2048.

Elliot, M., A. Galeotti, A. Koh, and W. Li (2021): "Market Segmentation through Information," Working Paper.

Ely, J. C. and J. Välimäki (2003): "Bad Reputation," *Quarterly Journal of Economics*, 118, 785–814.

Fainmesser, I. P. and A. Galeotti (2021): "The Market for Online Influence," *American Economic Journal: Microeconomics*, 13, 332–372.

Fudenberg, D., D. Kreps, and E. Maskin (1990): "Repeated Games with Long-run and Short-run Players," *Review of Economic Studies*, 57, 555–573.

Fundenberg, D. and D. K. Levine (1989): "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57, 759–778.

——— (1992): "Maintaining a Reputation when Strategies are Imperfectly Observed," *Review of Economic Studies*, 59, 561–579.

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

——— (2017): "Bayesian Persuasion with Multiple Senders and Rich Signal Spaces," *Games and Economic Behavior*, 104, 411–429.

Green, J. R. and N. Stokey (1978): "Two Representations of Information Structures and their Comparisons," Working Paper.

Haghpanah, N. and R. Siegel (2020): "Pareto Improving Segmentation of Multi-product Markets," Working Paper.

——— (forthcoming): "The Limits of Multi-Product Price Discrimination," *American Economic Review: Insight*.

Harbaugh, R. and E. Rasmusen (2018): "Coarse Grades: Informing the Public by Withholding Information," *American Economic Journal: Microeconomic*, 10, 210–235.

Hörner, J. (2002): "Reputation and Competition," *American Economic Review*, 92, 644–663.

Hörner, J. and N. Lambert (2021): "Motivational Ratings," *Review of Economic Studies*, 88, 1892–1935.

Johnson, J. P. and D. P. Myatt (2006): "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96, 756–784.

Kamenica, E. (2018): "Bayesian persuasion and information design," *Annual Review of Economics*.

Kim, K. and Y. Koh (2021): "Auctions with Flexible Information Acquisition," Working Paper.

Libgober, J. and X. Mu (2021): "Information Robustness in Intertemporal Pricing," *Review of Economic Studies*, 88, 1224–1252.

Liu, Q. (2011): "Information Acquisition and Reputation Dynamics," *Review and Economic Studies*, 78, 1400–1425.

Liu, Q. and A. Skrzypacz (2014): "Limited Records and Reputation Bubbles," *Journal of Economic Theory*, 151, 2–29.

Lizzeri, A. (1999): "Information Revelation and Certification Intermediaries," *RAND Journal of Economics*, 30, 214–231.

Mailath, G. J. and L. Samuelson (2001): "Who Wants a Good Reputation," *Review of Economic Studies*, 68, 415–441.

Maskin, E. and J. Tirole (2001): "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory*, 100, 191–219.

Mitchell, M. (2021): "Free Ad(vice): Internet Influencers and Disclosure Regulation," *RAND Journal of Economics*, 52, 3–21.

Ottaviani, M. and P. N. Sørensen (2006): "Reputational Cheap Talk," *RAND Journal of Economics*, 37, 441–466.

Pei, A. and D. Mayzlin (forthcoming): "Influencing the Influencers," *Marketing Science*.

Ravid, D., A.-K. Roseler, and B. Szentes (forthcoming): "Learning Before Trading: On the Inefficiency of Ignoring Free Information," *Journal of Political Economy*.

Roesler, A.-K. and B. Szentes (2017): "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107, 2072–2080.

Shi, X. (2012): "Optimal Auctions with Information Acquisition," *Games and Economic Behavior*, 74, 666–686.

Stahl, K. and R. Strausz (2017): "Certification and Market Transparency," *Review of Economic Studies*, 84, 1842–1868.

Terstiege, S. and C. Wasser (forthcoming): "Competitive Information Disclosure to an Auctioneer," *American Economic Journal: Microeconomics*.

Vellodi, N. (2021): "Ratings Design and Barriers to Entry," Working Paper.

Vong, A. (2021a): "Certifying Firms," Working Paper.

——— (2021b): "The Crisis of Expertise," Working Paper.

Yang, K. H. (2019): "Equivalence of Business Models for Informational Intermediaries," Working Paper.

——— (2021): "Efficient Market Demands in a Multi-Product Monopoly," *Journal of Economic Theory*, 197.

——— (forthcoming): "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and its Consequences," *American Economic Review*.

Appendix

This appendix contains proofs of results in Section 4. Proofs for other results are either more technical or about extensions, and thus are relegated to the online appendix.

A.1 Proof of Lemma 1

The "if" part immediately follows from the definition of strategies. For the "only if" part, it suffices to show that there is a profitable deviation only if there is a finite-shot deviation. To this end, consider any period t, any history h^t, and any strategy profile that gives the intermediary a finite continuation payoff. Let $\sigma|_{h^t}$ denote the intermediary's continuation strategy and let

$$\omega_t = \sum_{s=t}^{\infty} \delta^s m_s p_s q_s < \infty$$

denote the intermediary's continuation payoff, where $p_s \geq 0$, $D_s \in \mathcal{D}$ and $q_s \in [D_s(p_s^+), D_s(p_s)]$ are the price charged by the seller, the disclosure policy adopted by the intermediary and the tie-breaking rule chosen by the tie-breaker on path in each period $s \geq t$, respectively, and $\{m_s\}$ are the induced market sizes on path. Suppose that, when holding the seller's and the tie-breaker's strategies fixed, there is another continuation strategy $\tilde{\sigma}|_{h^t}$ at h^t that gives the intermediary a continuation payoff

$$\tilde{\omega}_t = \sum_{s=t}^{\infty} \delta^s \tilde{m}_s \tilde{p}_s \tilde{q}_s > \omega_t,$$

where $\tilde{p}_s \geq 0$, $\tilde{D}_s \in \mathcal{D}$, and $\tilde{q}_s \in [\tilde{D}_s(\tilde{p}_s^+), \tilde{D}_s(\tilde{p}_s)]$ are the price charged by the seller, the disclosure policy adopted by the intermediary, and the tie-breaking rule chosen by the tie-breaker in each period $s \geq t$ on the path induced by $\tilde{\sigma}|_{h^t}$, the seller's strategy and the tie-breaker's strategy, respectively, and $\{\tilde{m}_s\}$ are the associated market sizes. Furthermore, for any $T > t$, let $\tilde{\omega}_t^T$ be the intermediary's continuation payoff at history h^t when following $\tilde{\sigma}|_{h^t}$ until period $T > t$ and then return to $\sigma|_{h^t}$ from period T onward. Clearly, since stage game payoffs are bounded from below, both $\tilde{\omega}_t$ and $\tilde{\omega}_t^T$ are well-defined. Moreover, for $x \in \{\tilde{\omega}_t, \tilde{\omega}_t^T\}$, either $x = \infty$ or $x < \infty$. Clearly, if $\limsup_{T\to\infty} \tilde{\omega}_t^T = \infty$, then since $\omega_t < \infty$, there exists $\hat{T}$ such that $\tilde{\omega}_t^{\hat{T}} > \omega_t$ and hence deviating to $\tilde{\sigma}|_{h^t}$ for $\hat{T}$ periods then and return to $\sigma|_{h^t}$ is profitable for the intermediary. Thus, it is without loss to assume that $\limsup_{T\to\infty} \tilde{\omega}_t^T < \infty$. In the meantime, if $\tilde{\omega}_t = \infty$, then there exists $\hat{T}$ such that

$$\sum_{s=t}^{\hat{T}} \delta^s \tilde{m}_s \tilde{p}_s \tilde{q}_s > \omega_t,$$

which in turn implies that $\tilde{\omega}_t^{\hat{T}} > \omega_t$ since stage game payoffs are nonnegative. Therefore, it is also without loss to assume that $\tilde{\omega}_t < \infty$. Furthermore, since $\limsup_{T\to\infty} \tilde{\omega}_t^T < \infty$, $\{\tilde{\omega}_t^T\}$ is bounded. Thus, there exists a convergent subsequence $\{\tilde{\omega}_t^{T_n}\}$. We claim that it must be

$$\lim_{n\to\infty} \tilde{\omega}_t^{T_n} \geq \tilde{\omega}_t.$$

Indeed, suppose the contrary, that $\lim_{n\to\infty} \tilde{\omega}_t^{T_n} < \tilde{\omega}_t$. For any $n \in \mathbb{N}$, since $\tilde{\omega}_t^{T_n} < \infty$, it can be written as

$$\tilde{\omega}_t^{T_n} = \sum_{s=t}^{T_n} \delta^s \tilde{m}_s \tilde{p}_s \tilde{q}_s + \sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s^n \tilde{p}_s^n \tilde{q}_s^n,$$

where for any $n \in \mathbb{N}$, $\tilde{p}_s^n \geq 0$, $\tilde{D}_s^n \in \mathcal{D}$ and $\tilde{q}_s^n \in [\tilde{D}_s^n(\tilde{p}_s^{n+}), \tilde{D}_s^n(\tilde{p}_s^n)]$ are the price charged by the seller, the disclosure policies adopted by the intermediary, and the tie-breaking rule chosen by the tie-breaker in period $s \geq T_n + 1$ on path, respectively, and $\{\tilde{m}_s\}_{s=T_n+1}^{\infty}$ are the associated market sizes. Hence, for any $n \in \mathbb{N}$,

$$\tilde{\omega}_t^{T_n} - \tilde{\omega}_t = \sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s^n \tilde{p}_s^n \tilde{q}_s^n - \sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s \tilde{p}_s \tilde{q}_s.$$

Therefore, since $\tilde{\omega}_t < \infty$, and since $\sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s^n \tilde{p}_s^n \tilde{q}_s^n \leq \tilde{\omega}_t^{T_n}$ for all $n \in \mathbb{N}$,

$$\begin{aligned}
0 &> \lim_{n\to\infty} [\tilde{\omega}_t^{T_n} - \tilde{\omega}_t] \\
&= \lim_{n\to\infty} \left[\sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s^n \tilde{p}_s^n \tilde{q}_s^n - \sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s \tilde{p}_s \tilde{q}_s \right] \\
&= \lim_{n\to\infty} \sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s^n \tilde{p}_s^n \tilde{q}_s^n - \lim_{n\to\infty} \sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s \tilde{p}_s \tilde{q}_s \\
&= \lim_{n\to\infty} \sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s^n \tilde{p}_s^n \tilde{q}_s^n.
\end{aligned}$$

However, since stage game payoffs are nonnegative, $\delta^s \tilde{m}_s^n \tilde{p}_s^n \tilde{q}_s^n \geq 0$ for all $n \in \mathbb{N}$ and for all $s \geq T_n + 1$, which implies that

$$\lim_{n\to\infty} \sum_{s=T_n+1}^{\infty} \delta^s \tilde{m}_s^n \tilde{p}_s^n \tilde{q}_s^n \geq 0,$$

a contradiction. Thus, it must be that $\lim_{n\to\infty} \tilde{\omega}_t^{T_n} \geq \tilde{\omega}_t$.

As a result, since $\tilde{\omega}_t > \omega_t$, there exists $n \in \mathbb{N}$ such that $\tilde{\omega}_t^{T_n} > \omega_t$, as desired. ∎

A.2 Proof of Lemma 2

Consider any finite Markov perfect equilibrium. Since both the intermediary's and the seller's strategy do not depend on past histories in any finite Markov perfect equilibrium, the intermediary's normalized equilibrium continuation value in a given period must be a constant. Therefore, for any t, the intermediary's normalized continuation payoff at the beginning of period t can be written as $\omega^{\mathrm{M}} \in [0, \infty)$. Meanwhile, since the seller's strategy does not depend on history either, the price charged by the seller in period t must be a constant $p^{\mathrm{M}} \in [0, \infty)$ as well. Therefore, since both the intermediary and the seller are best responding in any finite Markov perfect equilibrium at any history, (4), (5), and (6) must hold.

Conversely, given any tuple $(\omega^{\mathrm{M}}, p^{\mathrm{M}}, \mathbf{D}^{\mathrm{M}})$ that satisfies the conditions required by the lemma, the strategy profile where the seller chooses p^{M} and the intermediary chooses $\mathbf{D}^{\mathrm{M}}(\cdot|p) \in \mathcal{D}$ whenever the seller chooses posted price $p \geq 0$ in the same period is immune to one-shot deviations. Moreover, since $\omega^{\mathrm{M}} < \infty$, Lemma 1 then implies that this strategy profile is indeed a subgame perfect equilibrium. This completes the proof. ∎

A.3 Proof of Lemma 3

First, notice that for any $D \in \mathcal{D}$,

$$(\mathbb{E}[v] - p)^+ \leq \int_p^\infty D(v)\,\mathrm{d}v \leq \int_p^\infty \overline{D}(v)\,\mathrm{d}v$$

for all $p \geq 0$. Moreover, notice that the function

$$\xi \mapsto \int_\xi^\infty \overline{D}(v)\,\mathrm{d}v - (p-\xi)\overline{D}(\xi)$$

is strictly decreasing on $[0,p]$, with a value of 0 at $\xi = v^{-1}(p)$ and a value of $\int_p^\infty \overline{D}(v)\,\mathrm{d}v$ at $\xi = p$, there must exist a unique $\xi(p) \in [v^{-1}(p), p]$ such that

$$\int_\xi^\infty \overline{D}(v)\,\mathrm{d}v - (p-\xi)\overline{D}(\xi) = \int_p^\infty D(v)\,\mathrm{d}v.$$

Now consider any $\widehat{D} \in \mathcal{D}$ such that

$$\int_{\xi(p)}^\infty \widehat{D}(v)\,\mathrm{d}v = \int_{\xi(p)}^\infty \overline{D}(v)\,\mathrm{d}v \tag{A.9}$$

and that

$$\widehat{D}(v) = \overline{D}(\xi(p)), \tag{A.10}$$

for all $v \in (\xi(p), p]$. Such $\widehat{D}$ exists since when $\widehat{D}$ is defined as

$$\widehat{D}(v) := \begin{cases} 1, & \text{if } v \in [0, \mathbb{E}[v|v \leq \xi(p)]] \\ \overline{D}(\xi(p)), & \text{if } v \in (\mathbb{E}[v|v \leq \xi(p)], \mathbb{E}[v|v \geq \xi(p)]] \\ 0, & \text{if } v \in (\mathbb{E}[v|v \geq \xi(p)], \infty) \end{cases},$$

we have

$$\int_{\xi(p)}^\infty \widehat{D}(v)\,\mathrm{d}v = \int_{\xi(p)}^\infty \overline{D}(v)\,\mathrm{d}v.$$

As a result, for any $\widehat{D}$ satisfying (A.9) and (A.10), by definition of $\xi(p)$,

$$\int_p^\infty \widehat{D}(v)\,\mathrm{d}v = \int_p^\infty D(v)\,\mathrm{d}v.$$

Moreover, $D \in \mathcal{D}$, (A.9) implies that

$$0 \leq \int_{\xi(p)}^\infty (\overline{D}(v) - D(v))\,\mathrm{d}v \leq \int_{\xi(p)}^p (\overline{D}(\xi(p)) - D(v))\,\mathrm{d}v \leq (p - \xi(p))(\overline{D}(\xi(p)) - D(p)),$$

and hence $D(p) \leq \overline{D}(\xi(p))$. Together, we have

$$\alpha p D(p) + \delta\omega\left(\gamma + \beta\int_p^\infty D(v)\,\mathrm{d}v\right) \leq \alpha p\overline{D}(\xi(p)) + \delta\omega\left(\gamma + \beta\left(\int_{\xi(p)}^\infty \overline{D}(v)\,\mathrm{d}v - (p-\xi(p))\overline{D}(\xi(p))\right)\right).$$

Lastly, notice that for any $\widehat{D}$ satisfying (A.9) and (A.10),

$$\alpha p\overline{D}(\xi(p)) + \delta\omega\left(\gamma + \beta\left(\int_{\xi(p)}^\infty \overline{D}(v)\,\mathrm{d}v - (p-\xi(p))\overline{D}(\xi(p))\right)\right) = \alpha p\widehat{D}(p) + \delta\omega\left(\gamma + \beta\int_p^\infty \widehat{D}(v)\,\mathrm{d}v\right).$$

As a result, for any D, there exists another $\widehat{D} \in \mathcal{D}$ satisfying (A.9) and (A.10) such that

$$\begin{aligned} \alpha p D(p) + \delta\omega\left(\gamma + \beta \int_p^\infty D(v)\,\mathrm{d}v\right) \leq & \alpha p \widehat{D}(p) + \delta\omega\left(\gamma + \beta \int_p^\infty \widehat{D}(v)\,\mathrm{d}v\right) \\ = & \alpha p \overline{D}(\xi(p)) + \delta\omega\left(\gamma + \beta\left(\int_{\xi(p)}^\infty \overline{D}(v)\,\mathrm{d}v - (p-\xi(p))\overline{D}(\xi(p))\right)\right). \end{aligned}$$

Therefore, the maximization problem

$$\sup_{D \in \mathcal{D}} \alpha p D(p) + \delta\omega\left(\gamma + \beta \int_p^\infty D(v)\,\mathrm{d}v\right)$$

can be simplified to

$$\max_{\xi \in [v^{-1}(p), p]} \alpha p \overline{D}(\xi) + \delta\omega\left(\gamma + \beta\left(\int_\xi^\infty \overline{D}(v)\,\mathrm{d}v - (p-\xi)\overline{D}(\xi)\right)\right), \tag{A.11}$$

which, by continuity of $\overline{D}$, has a solution. This implies that $\Delta(p|\omega)$ is nonempty. Moreover, the first-order Kuhn-Tucker condition of (A.11) implies its solution $\xi(p|\omega)$ is given by

$$\xi(p|\omega) = \max\left\{\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^+ p, v^{-1}(p)\right\}.$$

This in turn implies that any $\widehat{D} \in \mathcal{D}$ satisfying the condition given by the lemma must be in $\Delta(p|\omega)$. This completes the proof. ■

A.4 Proof of Lemma 4

By Lemma 3, for any selection $\mathbf{D}$ of $\Delta(\cdot|\omega)$,

$$\begin{aligned} p\mathbf{D}(p|p) = p\overline{D}(\xi(p|\omega)) = p\overline{D}\left(\max\left\{\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^+ p, v^{-1}(p)\right\}\right) \\ = \min\left\{p\overline{D}\left(\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^+ p\right), p\overline{D}(v^{-1}(p))\right\}, \end{aligned}$$

where the last equality follows from the fact that $\overline{D}$ is strictly decreasing. Furthermore, notice that for any p, if $1 \leq \alpha/\delta\beta\omega$, then

$$p\overline{D}\left(\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^+ p\right) = p.$$

Meanwhile, if $1 > \alpha/\delta\beta\omega$, let $\tilde{p} := (1 - \alpha/\delta\beta\omega)p$, then

$$p\overline{D}\left(\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^+ p\right) = \frac{\delta\beta\omega}{\delta\beta\omega - \alpha}\tilde{p}\overline{D}(\tilde{p}).$$

Thus, since $\overline{D}$ is regular, $p \mapsto \overline{D}((1 - \alpha/\delta\beta\omega)p)$ is quasi-concave as well. Lastly, by the definition of v^{-1} the function $p \mapsto \overline{D}(v^{-1}(p))$ is also quasi-concave. Together, the function

$$p \mapsto p\mathbf{D}(p|p)$$

is quasi-concave since it is a minimum of two quasi-concave functions. Thus, $\max_{p\geq 0} p\mathbf{D}(p|p)$ has a unique solution.

Moreover, if $1 \leq \alpha/\delta\beta\omega$, then $p\mathbf{D}(p|p) = p\overline{D}(v^{-1}(p))$ and hence $\tilde{p} = \mathbb{E}[v]$, which in turn implies that

$$\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^{+} \tilde{p} = 0 = v^{-1}(\tilde{p}) \leq \overline{p}.$$

Meanwhile, if $1 > \alpha/\delta\beta\omega$, notice that the function $p \mapsto \overline{D}(v^{-1}(p))$ is maximized at $p = \mathbb{E}[v]$ and that $\mathbb{E}[v]\overline{D}(v^{-1}(\mathbb{E}[v])) = \mathbb{E}[v] \geq \mathbb{E}[v]\overline{D}((1-\alpha/\delta\beta\omega)\mathbb{E}[v])$. As a result, since $p \mapsto p\mathbf{D}(p|p)$ is quasi-concave and hence single-peaked, it must attain its maximum at either price p such that $p\overline{D}(v^{-1}(p)) = p\overline{D}((1-\alpha/\delta\beta\omega)p)$ or the maximizer of $p\overline{D}((1-\alpha/\delta\beta\omega)p)$, whichever is smaller. Together with the fact that the maximizer of $p\overline{D}((1-\alpha/\delta\beta\omega)p)$ equals to

$$\operatorname*{argmax}_{\tilde{p}} \frac{\delta\beta\omega}{\delta\beta\omega - \alpha}\tilde{p}\overline{D}(\tilde{p}),$$

which is given by $\delta\beta\omega\overline{p}/(\delta\beta\omega - \alpha)$, it then follows that either

$$\left(1 - \frac{\alpha}{\delta\beta\omega}\right)\tilde{p} = v^{-1}(\tilde{p}) \text{ and } \tilde{p} \leq \frac{\delta\beta\omega}{\delta\beta\omega - \alpha}\overline{p}$$

or

$$\tilde{p} = \frac{\delta\beta\omega}{\delta\beta\omega - \alpha}\overline{p} \leq v(\overline{p}).$$

As a result, it must be that

$$v^{-1}(\tilde{p}) \leq \left(1 - \frac{\alpha}{\delta\beta\omega}\right)^{+} \tilde{p} \leq \overline{p},$$

with at least one inequality binding.

Lastly, suppose that

$$\int_{\tilde{p}}^{\infty} \mathbf{D}(v|\tilde{p})\,\mathrm{d}v = 0.$$

Then, by Lemma 3,

$$\int_{\xi(\tilde{p}|\omega)}^{\infty} \overline{D}(v)\,\mathrm{d}v - (\tilde{p} - \xi(\tilde{p}|\omega))\overline{D}(\xi(\tilde{p}|\omega)) = 0,$$

which is equivalent to

$$\mathbb{E}[v|v \geq \xi(\tilde{p}|\omega)] = \tilde{p} \iff \xi(\tilde{p}|\omega) = v^{-1}(\tilde{p}).$$

Moreover, notice that for any $p \in [0, \mathbb{E}[v]]$

$$p\overline{D}(v^{-1}(p)) = p \geq p\overline{D}\left(\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^{+} p\right)$$

and that $p \mapsto p\overline{D}(v^{-1}(p))$ is uniquely maximized at $p = \mathbb{E}[v]$, single-peakness of $p \mapsto \mathbf{D}(p|p)$ and $\xi(\tilde{p}|\omega) = v^{-1}(\tilde{p})$ then implies that

$$\tilde{p}\overline{D}(v^{-1}(\tilde{p})) = \tilde{p}\overline{D}\left(\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^{+} \tilde{p}\right),$$

which in turn implies that

$$\left(1 - \frac{\alpha}{\delta\beta\omega}\right)^{+} \tilde{p} = v^{-1}(\tilde{p}).$$

Conversely, suppose that

$$\left(1-\frac{\alpha}{\delta\beta\omega}\right)^{+}\tilde{p}=v^{-1}(\tilde{p}).$$

Then $\xi(\tilde{p}|\omega)=v^{-1}(\tilde{p})$ and hence, by Lemma 3,

$$\int_{\tilde{p}}^{\infty}\mathbf{D}(v|\tilde{p})\,\mathrm{d}v=\int_{v^{-1}(\tilde{p})}^{\infty}\overline{D}(v)\,\mathrm{d}v-(\tilde{p}-v^{-1}(\tilde{p}))\overline{D}(v^{-1}(\tilde{p}))=\overline{D}(v^{-1}(\tilde{p}))(\tilde{p}-v(v^{-1}(\tilde{p})))=0.$$

This completes the proof. ■

A.5 Proof of Theorem 1

We first show that any $(r^{\mathrm{M}},p^{\mathrm{M}},\sigma^{\mathrm{M}},\omega^{\mathrm{M}},\{m_t^{\mathrm{M}}\})$ described in the statement of the theorem is indeed a finite Markov perfect equilibrium. To this end, we will show that for any such tuple, there exists $\mathbf{D}^{\mathrm{M}}:\mathbb{R}_+\to\mathcal{D}$ such that $(\omega^{\mathrm{M}},p^{\mathrm{M}},\mathbf{D}^{\mathrm{M}})$ satisfies the conditions of Lemma 2. Consider three cases separately.

Case 1: $\beta\in[0,\underline{\beta}]$.
In this case, notice that

$$\delta\beta\omega^{\mathrm{M}}=\delta\beta\frac{\alpha\mathbb{E}[v]}{1-\gamma\delta}=\frac{\beta\alpha}{\underline{\beta}}\leq\alpha,$$

and therefore $\omega^{\mathrm{M}}\leq\alpha/\delta\beta$. Consider any selection $\mathbf{D}^{\mathrm{M}}$ of $\Delta(\cdot|\omega^{\mathrm{M}})$. Since $p^{\mathrm{M}}=\mathbb{E}[v]$ and thus $v^{-1}(p^{\mathrm{M}})=0$, Lemma 4 implies that $p^{\mathrm{M}}\in\operatorname{argmax}_p p\mathbf{D}^{\mathrm{M}}(p|p)$, which establishes (5). Meanwhile, by Lemma 4, it must be that

$$\int_{p^{\mathrm{M}}}^{\infty}\mathbf{D}^{\mathrm{M}}(v|p^{\mathrm{M}})\,\mathrm{d}v=0.$$

Moreover, given $p^{\mathrm{M}}=\mathbb{E}[v]$, Lemma 3 implies that $\mathbf{D}^{\mathrm{M}}(v|p^{\mathrm{M}})=\overline{D}(0)=1$. Together,

$$\begin{aligned}\sup_{D\in\mathcal{D}}\left[\alpha p^{\mathrm{M}}D(p^{\mathrm{M}})+\delta\left(\gamma+\beta\int_{p^{\mathrm{M}}}^{\infty}D(v)\,\mathrm{d}v\right)\omega^{\mathrm{M}}\right]=&\alpha p^{\mathrm{M}}\mathbf{D}^{\mathrm{M}}(p^{\mathrm{M}}|p^{\mathrm{M}})+\delta\left(\gamma+\beta\int_{p^{\mathrm{M}}}^{\infty}\mathbf{D}^{\mathrm{M}}(v|p^{\mathrm{M}})\,\mathrm{d}v\right)\\=&\alpha\mathbb{E}[v]+\gamma\delta\omega^{\mathrm{M}}\\=&\omega^{\mathrm{M}}\\=&\frac{\alpha p^{\mathrm{M}}\mathbf{D}^{\mathrm{M}}(p^{\mathrm{M}}|p^{\mathrm{M}})}{1-\delta\left(\gamma+\beta\int_{p^{\mathrm{M}}}^{\infty}\mathbf{D}^{\mathrm{M}}(v|p)\,\mathrm{d}v\right)}\end{aligned}$$

which establishes (4) and (6).

Case 2: $\beta\in(\underline{\beta},\overline{\beta})$.
In this case, take any selection $\mathbf{D}^{\mathrm{M}}$ of $\Delta(\cdot|\omega^{\mathrm{M}})$. Notice that by definition, $1>\alpha/\delta\beta\omega^{\mathrm{M}}$, and hence

$$\left(1-\frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)^{+}p^{\mathrm{M}}=\left(1-\frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)p^{\mathrm{M}}=v^{-1}(p^{\mathrm{M}}),$$

which in turn implies that, by Lemma 4,

$$\int_{p^{\mathrm{M}}}^{\infty}\mathbf{D}^{\mathrm{M}}(v|p^{\mathrm{M}})\,\mathrm{d}v=0$$

and therefore,

$$\omega^{\mathrm{M}} = \alpha p^{\mathrm{M}}\overline{D}(\xi(p^{\mathrm{M}}|\omega^{\mathrm{M}})) + \gamma\delta\omega^{\mathrm{M}} = \alpha p^{\mathrm{M}}\mathbf{D}^{\mathrm{M}}(p^{\mathrm{M}}|p^{\mathrm{M}}) + \delta\left(\gamma + \beta\int_{p^{\mathrm{M}}}^{\infty}\mathbf{D}^{\mathrm{M}}(v|p^{\mathrm{M}})\,\mathrm{d}v\right)\omega^{\mathrm{M}},$$

which establishes (4). Furthermore, since $p^{\beta} < \overline{p}$, $p^{\mathrm{M}} < \delta\beta\omega^{\mathrm{M}}\overline{p}/(\delta\beta\omega^{\mathrm{M}} - \alpha)$ and hence p^{M} is the unique maximizer of $p\overline{D}(\xi(p|\omega^{\mathrm{M}}))$ according to Lemma 4. Thus, by Lemma 3, $(\omega^{\mathrm{M}}, p^{\mathrm{M}}, \mathbf{D}^{\mathrm{M}})$ satisfies (5) and (6).

Case 3: $\beta = \overline{\beta}$.
In this case, consider any selection $\mathbf{D}^{\mathrm{M}}$ of $\Delta(\cdot|\omega^{\mathrm{M}})$. By definition, $1 > \alpha/\delta\beta\omega^{\mathrm{M}}$ and

$$\left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)p^{\mathrm{M}} > v^{-1}(p^{\mathrm{M}}),$$

and thus, by Lemma 3 and Lemma 4,

$$p^{\mathrm{M}}\mathbf{D}^{\mathrm{M}}(p^{\mathrm{M}}|p^{\mathrm{M}}) = p^{\mathrm{M}}\overline{D}\left(\left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)p^{\mathrm{M}}\right) \geq p\overline{D}\left(\max\left\{\left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)p, v^{-1}(p)\right\}\right) = p\mathbf{D}^{\mathrm{M}}(p|p)$$

As a result, $(\omega^{\mathrm{M}}, p^{\mathrm{M}}, \mathbf{D}^{\mathrm{M}})$ satisfies (4), (5), and (6) and therefore induces a stationary equilibrium, as desired.

We now show that for any finite Markov perfect equilibrium, its outcome $(r^{\mathrm{M}}, p^{\mathrm{M}}, \sigma^{\mathrm{M}}, \omega^{\mathrm{M}}, \{m_t^{\mathrm{M}}\})$ must satisfy the conditions given by Theorem 1. By Lemma 2, there exists $(\omega^{\mathrm{M}}, p^{\mathrm{M}}, \mathbf{D}^{\mathrm{M}})$ satisfying (4), (5), and (6) such that $r^{\mathrm{M}} = p^{\mathrm{M}}\mathbf{D}^{\mathrm{M}}(p^{\mathrm{M}}|p^{\mathrm{M}})$, $\sigma^{\mathrm{M}} = \int_{p^{\mathrm{M}}}^{\infty}\mathbf{D}^{\mathrm{M}}(v|p^{\mathrm{M}})\,\mathrm{d}v$ and $m_t^{\mathrm{M}} = (1 + \beta\sigma^{\mathrm{M}})^t$. It follows immediately that $r^{\mathrm{M}}, \sigma^{\mathrm{M}}, \{m_t^{\mathrm{M}}\}$ satisfy the condition given by Theorem 1 if ω^{M} and p^{M} satisfy these conditions. Thus, it suffices to show that $\omega^{\mathrm{M}}, p^{\mathrm{M}}$ satisfy these conditions. To this end, notice that By Lemma 4,

$$v^{-1}(p^{\mathrm{M}}) \leq \left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)^{+} p^{\mathrm{M}} \leq \overline{p}, \tag{A.12}$$

with at least one inequality binding. Now consider three cases separately.

Case 1: $\omega^{\mathrm{M}} \leq \alpha/\delta\beta$.
In this case, it immediately follows that

$$\left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)^{+} p^{\mathrm{M}} = 0 = v^{-1}(p^{\mathrm{M}})$$

and hence $p^{\mathrm{M}} = \mathbb{E}[v]$, which in turn, by (4), implies that $\omega^{\mathrm{M}} = \alpha\mathbb{E}[v]/(1 - \gamma\delta)$. For this to be consistent with $\omega^{\mathrm{M}} \leq \alpha/\delta\beta$, it must be that $\beta \leq \underline{\beta}$.

Case 2: $\omega^{\mathrm{M}} > \alpha/\delta\beta$ and

$$\left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)p^{\mathrm{M}} = v^{-1}(p^{\mathrm{M}}). \tag{A.13}$$

In this case, Lemma 3 implies that

$$\omega^{\mathrm{M}} = \delta\left(\gamma + \beta\int_{\left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right)p^{\mathrm{M}}}\overline{D}(v)\,\mathrm{d}v\right)\omega^{\mathrm{M}},$$

and hence, together with (A.12), it must be that $\beta \in [\underline{\beta}, \overline{\beta}]$ and

$$\left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right) p^{\mathrm{M}} = p^{\beta}.$$

Meanwhile, since (A.13) is equivalent to

$$\omega^{\mathrm{M}} = g^{\beta}\left(\left(1 - \frac{\alpha}{\delta\beta\omega^{\mathrm{M}}}\right) p^{\mathrm{M}}\right),$$

it must be that $\omega^{\mathrm{M}} = g^{\beta}(p^{\beta})$ and hence $p^{\mathrm{M}} = v(p^{\beta})$.

Case 3: $\omega^{\mathrm{M}} > \alpha/\delta\beta$ and

$$p^{\mathrm{M}} = \frac{\delta\beta\omega^{\mathrm{M}}}{\delta\beta\omega^{\mathrm{M}} - \alpha}\overline{p}. \tag{A.14}$$

In this case, Lemma 3 implies that

$$\omega^{\mathrm{M}} = \delta\left(\gamma + \beta\int_{\overline{p}}^{\infty} \overline{D}(v)\,\mathrm{d}v\right)\omega^{\mathrm{M}},$$

which means this case can only occur when $\beta = \overline{\beta}$.

Together with observations that $p^{\beta} = 0$ if and only if $\beta \leq \underline{\beta}$, that $p^{\beta} = \overline{p}$ if and only if $\beta = \overline{p}$, and that the second inequality of (A.12) is equivalent to $\omega^{\mathrm{M}} \geq g^{\beta}(\overline{p})$, it then follows that $\omega^{\mathrm{M}}, p^{\mathrm{M}}$ must be the same as described in Theorem 1 in all three cases. This completes the proof. ■

A.6 Proof of Theorem 2

We first show that any Markov perfect equilibrium must be finite whenever $\beta < \widetilde{\beta}$. Suppose that $\mathbf{z}^{\mathrm{M}} = (r^{\mathrm{M}}, \sigma^{\mathrm{M}}, \omega^{\mathrm{M}}, p^{\mathrm{M}}, \{m_t^{\mathrm{M}}\})$ is a Markov perfect equilibrium outcome, and suppose that, by way of contradiction, $\omega^{\mathrm{M}} = \infty$. Then it must be that

$$\delta(\gamma + \beta\sigma^{\mathrm{M}}) \geq 1,$$

which in turn implies that

$$\sigma^{\mathrm{M}} \geq \int_{p^{\beta}}^{\infty} \overline{D}(v)\,\mathrm{d}v.$$

Meanwhile, since the total surplus is at most $\mathbb{E}[v]$ and since the seller's revenue r^{M} must be at least r^* in any subgame perfect equilibrium, it must be that

$$\sigma^{\mathrm{M}} \leq \mathbb{E}[v] - r^*.$$

Together, we have

$$\int_{p^{\beta}}^{\infty} \overline{D}(v)\,\mathrm{d}v \leq \mathbb{E}[v] - r^*,$$

which is equivalent to $\beta \geq \tilde{\beta}$, a contradiction.

Now suppose that $\beta \geq \widetilde{\beta}$. We first show that any infinite Markov perfect equilibrium outcome $\mathbf{z}^{\mathrm{M}} = (r^{\mathrm{M}}, \sigma^{\mathrm{M}}, \omega^{\mathrm{M}}, p^{\mathrm{M}}, \{m_t^{\mathrm{M}}\})$ must satisfy condition 2 of the theorem. Indeed, by definition, $\omega^{\mathrm{M}} = \infty$ and $m_t = (\gamma + \beta\sigma^{\mathrm{M}})^t$ for all t, which in turn implies that $\delta(\gamma + \beta\sigma^{\mathrm{M}}) \geq 1$. Rearranging, we have

$$\sigma^{\mathrm{M}} \geq \frac{1 - \gamma\delta}{\delta\beta}.$$

Meanwhile, since the seller's revenue in every period must be at least r^* in any subgame perfect equilibrium, $r^{\mathrm{M}} \geq r^*$. Therefore, since the efficient surplus is $\mathbb{E}[v]$, we have

$$\sigma^{\mathrm{M}} \leq \mathbb{E}[v] - r^{\mathrm{M}} \leq \mathbb{E}[v] - r^*.$$

Furthermore, notice that since $\sigma^{\mathrm{M}} = \int_{p^{\mathrm{M}}}^{\infty} D(v)\,\mathrm{d}v$ for some $D \in \mathcal{D}$, it must be that $\int_{p^{\mathrm{M}}}^{\infty} \overline{D}(v)\,\mathrm{d}v \geq \sigma^{\mathrm{M}}$. Given any σ and any p such that $\int_{p}^{\infty} \overline{D}(v)\,\mathrm{d}v \geq \sigma$, notice that the function

$$q \mapsto S(q) - pq$$

is quasi-concave and hence the equation

$$S(q) - pq = \sigma$$

has at most two solutions, denoted as $\underline{q}(p,\sigma) \leq \overline{q}(p,\sigma)$. It then follows that

$$\underline{q}(p,\sigma) \leq D(p) \leq \overline{q}(p,\sigma)$$

for all $D \in \mathcal{D}$ such that $\int_{p}^{\infty} D(v)\,\mathrm{d}v = \sigma$. Moreover, since $q \mapsto S(q) - pq$ is increasing, $\underline{q}(\cdot,\sigma)$ is decreasing in p. As a result, for any $D \in \mathcal{D}$ and for any $p \geq 0$ such that $\int_{p}^{\infty} D(v)\,\mathrm{d}v = \sigma$,

$$pD(p) \geq p\underline{q}(p,\sigma) \geq (\mathbb{E}[v] - \sigma)\underline{q}(\mathbb{E}[v] - \sigma, \sigma) = S(\underline{q}(\mathbb{E}[v] - \sigma, \sigma)) - \sigma.$$

Lastly, notice that by the definition of $\underline{q}(p,\sigma)$ and $\zeta(p)$, we have $\underline{q}(\mathbb{E}[v] - \sigma, \sigma) = \overline{D}(\zeta(\mathbb{E}[v] - \sigma))$. Together, it must be that

$$\max\left\{S(\overline{D}(\zeta(\mathbb{E}[v] - \sigma^{\mathrm{M}}))) - \sigma^{\mathrm{M}}, r^*\right\} \leq r^{\mathrm{M}} \leq \mathbb{E}[v] - \sigma^{\mathrm{M}}.$$

In the meantime, since $\int_{p^{\mathrm{M}}}^{\infty} D(v)\,\mathrm{d} = \sigma^{\mathrm{M}}$ and $p^{\mathrm{M}}D(p^{\mathrm{M}}) = r^{\mathrm{M}}$ for some $D \in \mathcal{D}$, it must be that $p^{\mathrm{M}} \geq \mathbb{E}[v] - \sigma^{\mathrm{M}}$ and that

$$S\left(\frac{r^{\mathrm{M}}}{p^{\mathrm{M}}}\right) - r^{\mathrm{M}} \geq \sigma^{\mathrm{M}},$$

which is equivalent to

$$\mathbb{E}[v] - \sigma^{\mathrm{M}} \leq p^{\mathrm{M}} \leq \frac{r^{\mathrm{M}}}{S^{-1}(r^{\mathrm{M}} + \sigma^{\mathrm{M}})},$$

as desired.

Conversely, suppose that $\mathrm{z}^{\mathrm{M}} = (r^{\mathrm{M}}, \sigma^{\mathrm{M}}, \omega^{\mathrm{M}}, p^{\mathrm{M}}, \{m_t^{\mathrm{M}}\})$ satisfies condition 2 of the theorem. We now construct a Markov perfect equilibrium whose outcome is $\mathbf{z}^{\mathrm{M}}$. To this end, let the seller's strategy be p^{M} for all periods and for all histories. Consider a strategy for the intermediary as follows: For any history, if the seller charges $p \neq p^{\mathrm{M}}$ in the same period, choose a solution of

$$\min_{D \in \mathcal{D}} pD(p^+).$$

Otherwise, if the seller charges p^{M} choose $D \in \mathcal{D}$ such that

$$D(p^{\mathrm{M}}) = \frac{r^{\mathrm{M}}}{p^{\mathrm{M}}}$$

and that

$$\int_{p^{\mathrm{M}}}^{\infty} D(v)\,\mathrm{d}v = \sigma^{\mathrm{M}}.$$

Using the same arguments as above, since

$$\begin{aligned}\frac{1-\gamma\delta}{\delta\beta} \leq& \sigma^{\mathrm{M}} \leq \mathbb{E}[v] - r^* \\ \mathbb{E}[v] - \sigma^{\mathrm{M}} \leq& p^{\mathrm{M}} \leq \frac{r^{\mathrm{M}}}{S^{-1}(r^{\mathrm{M}} + \sigma^{\mathrm{M}})}, \\ \max\left\{S(\overline{D}(\zeta(\mathbb{E}[v] - \sigma^{\mathrm{M}}))) - \sigma^{\mathrm{M}}, r^*\right\} \leq& r^{\mathrm{M}} \leq \mathbb{E}[v] - \sigma^{\mathrm{M}},\end{aligned}$$

such $D \in \mathcal{D}$ exists. Finally, consider the following strategy for the tie-breaker: For any history, if the seller chooses $p \neq p^{\mathrm{M}}$ and the intermediary chooses any $D \in \mathcal{D}$ in the same period, then chooses $q = D(p^+)$. If the seller chooses p^{M} and the intermediary chooses any $D \in \mathcal{D}$ in the same period, then choose $q = D(p)$.

By construction, the above strategy profile is Markov. Moreover, since $r^{\mathrm{M}} \geq r^*$ and since the seller can get at most r^* if he deviates given the intermediary's and the tie-breaker's strategies, the seller would never deviate. As for the intermediary, when the price is p^{M} and the market size is m, by construction, the market size in the next period is $(\gamma + \beta\sigma^{\mathrm{M}})m$. Since $\sigma^{\mathrm{M}} \geq (1-\gamma\delta)/\delta\beta$, the intermediary's payoff is $\omega^{\mathrm{M}} = \infty$ given the seller always charges p^{M} and given the tie-breaker's strategy. Meanwhile, if the seller charges $p \neq p^{\mathrm{M}}$, then given that the seller charges p^{M} in all future periods, choosing the solution of

$$\min_{D \in \mathcal{D}} pD(p^+)$$

and then return to the aforementioned strategy in future periods still gives a present discounted profit of ∞. Together, the intermediary would not deviate given the seller's and the tie-breaker's strategies. Thus, this strategy profile is indeed a Markov perfect equilibrium with outcome $\mathbf{z}^{\mathrm{M}}$.

Finally, since

$$\frac{1-\gamma\delta}{\delta\beta} \leq \mathbb{E}[v] - r^*$$

there exists $\mathbf{z}^{\mathrm{M}}$ that satisfies condition 2 of the theorem whenever $\beta \geq \widetilde{\beta}$. Together with the proofs above, it then follows that there exists an infinite Markov perfect equilibrium whenever $\beta \geq \widetilde{\beta}$. This completes the proof. ■