

Should platforms be allowed to sell on their own marketplaces?*

Andrei Hagiu[†]

Tat-How Teh[‡]

Julian Wright[§]

July 30, 2021

Abstract

A growing number of platforms including those run by Amazon, Apple, Google and Microsoft operate in a dual mode: running marketplaces for third-party products, while also selling their own products on those marketplaces. We build a model to explore the implications of this controversial practice, taking into account the endogeneity of third-party sellers' innovation decisions. We analyze the tradeoffs that arise from a ban on the dual mode, showing how such a ban can harm consumer surplus and welfare even when the platform would otherwise engage in product imitation and self-preferencing. Instead, in the empirically most relevant scenarios, policies that prevent platform imitation and self-preferencing lead to better outcomes than an outright ban on the dual mode.

1 Introduction

An increasing number of e-commerce players such as Amazon, JD.com, Target, and Walmart, are acting both as marketplaces, i.e. enabling third party sellers to sell to consumers, and as sellers, i.e. selling products under their own name. Similarly, Apple's App Store, Amazon's AWS Marketplace, Google's Chrome Web Store, Cloud Marketplace and Play Store, Intuit's Quickbooks App Store, Microsoft's Azure Marketplace, Windows Games Store and Xbox Games, Nintendo's Game Store, Salesforce's AppExchange, Shopify's App Store, Sony's PlayStation Store, all sell their own apps/games alongside third-party apps/games on their marketplaces.

This practice has raised regulatory concerns over the lack of a level playing field, and has led to investigations and new proposed legislation in many jurisdictions. In February 2019, India introduced new laws to force the separation of the two types of businesses, leading Amazon and the Walmart-backed Flipkart to change their business practices there so they stopped selling under their own name. The Ending Platform Monopolies Act proposed in the U.S. in 2021 includes provisions aimed at stopping "Big Tech" firms from selling their own competing products or apps in competition with third-parties on their respective marketplaces.

*We thank Heski Bar-Isaac, Justin Johnson, Jan Krämer, Martin Peitz, Tommaso Valletti, as well as other participants in talks at APIOC 2019, CEPR VIOS 2020, the 18th ZEW Conference on Information and Communication Technologies, University of Manchester, Lingnan University, National University of Singapore, Technology and Declining Economic Dynamism Conference at Boston University, Wharton Technology and Innovation Conference 2021, for their helpful comments. We gratefully acknowledge research funding from the Singapore Ministry of Education Social Science Research Thematic Grant, MOE2017-SSRTG-023. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of the Singapore Ministry of Education or the Singapore Government.

[†]Boston University Questrom School of Business. E-mail: ahagiu@bu.edu

[‡]School of Management and Economics and Shenzhen Finance Institute, The Chinese University of Hong Kong, Shenzhen, China, e-mail: tehtathow@cuhk.edu.cn

[§]Department of Economics, National University of Singapore, E-mail: jwright@nus.edu.sg

In this paper we build a tractable model of a platform that can adopt a dual mode, in which it sells products in its own name (i.e. seller mode) alongside third-party sellers who sell competing products (i.e. marketplace mode) to explore the welfare implications of this practice. Specifically, we use the model to study how the platform’s optimal choice of mode changes when the dual mode is outlawed, and derive the implications for consumer surplus and total welfare of such a ban. We also conduct a similar analysis for several alternative policy options.

There are many reasons why it might be efficient (and pro-competitive) for the same platform to act as a seller for some types of products and as a marketplace for others. Most obviously, some types of products may be more efficiently provided by the platform, while others may be more efficiently provided by third-party sellers, and offering all of them in the same place provides one-stop shopping benefits to consumers. This means a blanket ban on the dual mode, i.e. one that requires platforms to choose the same mode (either seller or marketplace) across all products, is more likely to be harmful for consumers and welfare than just banning the dual mode at the level of an individual product or a narrowly defined product category. For this reason, in this paper we study the welfare implications of banning the dual mode for an individual product category. If we find a narrow ban on the dual mode (at the product level) is bad for consumers or welfare, then a wider ban that prevents the platform acting as a seller for some products and a marketplace on others is even more likely to be harmful.

A number of antitrust concerns have been raised when a dominant platform adopts the dual mode at the product level. These center on the possibility that the platform may want to favor the products it sells and so distort competition in the marketplace, leading to unfair competition. This can happen in at least two important ways.¹ One is that the platform obtains proprietary information on the third-party sellers’ products (e.g. detailed demand and pricing data, data on users’ search behavior, what items they return and their reasons for doing so) via its marketplace, and then uses that opportunistically to decide whether to copy and compete on the more successful offerings, potentially leading to reduced incentives for third-party sellers to invest or innovate.² A second concern is that the platform can steer consumers towards its own offerings rather than those offered by third-party sellers by displaying its own offerings more prominently, a practice that has become known as self-preferencing. For example, Amazon can do this through its Buy Box, which around 85% of consumers click on to complete their order. This allocates a seller to the consumer according to a secret algorithm that Amazon controls, and oftentimes the allocated seller is Amazon itself.

To model these practices, we assume that there is a platform M that can function as a seller and/or a marketplace, a fringe of small third-party sellers that all sell an identical product, and an innovative seller S that has a better product in the same category as the fringe sellers, and indeed can invest more in making its product even better. The platform M also offers a product in the same category which can be better or worse than the one offered by the fringe sellers (or equivalently, which could involve a lower or a higher cost of production). While consumers are aware of fringe sellers and can bypass the platform to buy directly from them if they prefer, they rely on M to discover the innovative seller’s product. This captures that in practice many consumers rely on large marketplaces like Amazon’s to find out about the existence of a seller offering a new product variant.³ In the baseline version of the model without the possibility of self-preferencing, consumers become aware of S ’s existence as long as S lists on M ’s marketplace. Once they become aware of S , they also learn its price in the direct channel and are free to switch to buy from S directly if they prefer (i.e. showrooming is possible).

¹Both of these concerns are recognized in the recently proposed Digital Markets Act in Europe and the American Choice and Innovation Online Act in the U.S.

²See Mattioli (2020) for reports that employees working as part of Amazon’s private labels group used data from its marketplace to inform their decisions about launching private label products.

³Note that in the U.S., significantly more consumers start their product searches on Amazon than on Google search (Statista, July 7, 2021).

Apart from potentially discovering S if they go to M , consumers also obtain some convenience benefit from using the platform to complete transactions. While we allow the different channels (on-platform or direct) and the different sellers (M , S , or fringe sellers) to offer different values to consumers, consumers are all identical in their preferences except that they differ in the value of their outside option which is drawn randomly from a general distribution. This setup preserves tractability while allowing for the level of prices to matter for aggregate demand and efficiency.

We explore three different business models for the platform: marketplace mode (only facilitating transactions by third-party sellers who set their own prices for these transactions), seller mode (only selling its own product in competition with outside sellers), and dual mode (operating in both modes, selling in its own name at the same time as facilitating transactions by third-party sellers). We start with a baseline setting in which imitation and self-preferencing are not possible.

Our first finding is that the dual mode is always better for the platform than just being a marketplace. Even though the possibility of showrooming means M chooses the same commission in both cases (i.e. S would induce consumers to buy directly if M 's commission is too high), the competitive pressure of M 's offering lowers S 's price in equilibrium, leading to more transactions on the marketplace and more commission revenue for M . The benefits of this price squeeze is also the reason the dual mode can be better for M than just being a seller. But relative to just being a seller, by choosing the dual mode (and without the possibility to steer), M essentially commits to let all consumers become aware of S 's product. The resulting competition hurts M when its own offering is valued more highly than the fringe sellers' products since it can no longer fully capture the additional value it offers. Thus, M prefers the dual mode over the seller mode provided its own offering is not too much better than the offerings of the fringe sellers.

We then analyze the effect of a ban on the dual mode, taking into account that M endogenously decides which mode to switch to in response to the ban. A ban on the dual mode, whenever it is relevant, always results in lower consumer surplus. The ban also decreases total welfare whenever M 's product creates more value for consumers (or can be produced at a lower cost) than the one provided by fringe sellers, which we think is the empirically more relevant case for the types of platforms that regulators are considering imposing such bans on. In this case, M would switch to seller mode in response to the ban, resulting in a decrease in consumer surplus due to the weaker price competition that arises when consumers are unaware of S 's existence. The corresponding decrease in welfare reflects two sources of inefficiencies in the seller mode: (i) there are fewer transactions; and (ii) consumers can no longer combine S 's superior product with M 's convenience benefit. Even when the ban on the dual mode leads M to switch to marketplace mode, consumer surplus is still lower (due to the absence of the price squeeze which puts competitive pressure on S 's price), but the welfare effect can go in either direction given the possibility that S may invest less when M operates in dual mode than when it operates in marketplace mode. Surprisingly though, S may sometimes invest more in dual mode. This possibility arises because the price squeeze in dual mode increases transaction volume, which can more than offset the lower margins that S obtains, thus increasing S 's marginal gain from innovation.

We then use our framework to explore the practices of product imitation and self-preferencing that have raised antitrust scrutiny. To do so, we modify our baseline model by assuming that whenever S is hosted on the marketplace: (i) M can perfectly and costlessly imitate S 's innovative product; and (ii) M can choose whether to disclose the availability of S on its marketplace. Thus, we allow for the most extreme forms of imitation and self-preferencing in dual mode, which maximizes the potential harm (to innovation, consumer surplus and welfare) from the dual mode. There are three key implications from these modifications. First, the dual mode is now always weakly better than either pure mode from M 's perspective. Second, anticipating that M will perfectly imitate its product, S has no incentive to invest in improving its product when M operates in dual mode. And third, M 's ability to perfectly steer

consumers away from S 's product essentially eliminates any on-platform price competition in dual mode, while at the same time raising the commission that M can charge.

We consider four possible policy interventions to address the harm arising from product imitation and self-preferencing in dual mode.

First, we reconsider banning the dual mode. This ban turns out to no longer matter for consumers because the weaker on-platform price competition (as a result of self-preferencing) means that consumers do not benefit from the existence of S in dual mode (which mirrors the situation in either the marketplace or seller modes). Nonetheless, to the extent it has any effect, banning the dual mode still has a negative effect on welfare in the empirically more relevant case that M has an advantage in selling compared to fringe sellers. One of the channels through which the ban on the dual mode lowers welfare (fewer transactions due to higher prices) is now shut down, but the other remains (consumers cannot combine S 's superior product with M 's convenience benefit).

Second, we show that banning product imitation alone either has no effect or increases consumer surplus and welfare. Interestingly, the latter happens when M also benefits from the ban: this is because the ban allows M to credibly commit not to imitate, which benefits M (as well as consumers and welfare) when it induces a sufficiently large innovation from S .

Third, we show that banning self-preferencing alone has more ambiguous implications. In the case that M continues in dual mode after the ban, the ban restores on-platform competition and the showrooming constraint on M 's commission. Both of these decrease the final price level and benefit consumers (as well as total welfare, via the number of transactions). Nonetheless, because imitation is still possible, S 's innovation level remains at its minimum. And in other cases, after the ban, M may switch to seller mode, which as noted above, lowers welfare whenever there is any effect at all, for the same reason as above.

Finally, we consider banning both product imitation and self-preferencing, which can result in M either choosing the seller mode or continuing in dual mode. In the latter case, consumer surplus and welfare increase because banning imitation and self-preferencing restores the showrooming constraint, so S 's innovation incentive is back to what it was in the baseline dual mode without these practices. Interestingly, we find that banning imitation and banning self-preferencing have a synergistic effect, in the sense that the condition for innovation to increase following the ban is more likely to hold compared to if we just ban imitation alone. Nevertheless, in case the ban leads to M adopting the seller mode, welfare can be lower for the same reason as before.

These results reflect that banning product imitation and self-preferencing address the negative consequences of each of these practices, while at the same time preserving some of the benefits of the dual mode we found in the baseline setting. Comparing the implications of all four policy interventions, our results suggest that a structural ban on the dual mode is a less effective intervention than behavioural remedies. The only exception to this, where things become ambiguous, is the empirically less relevant case in which the fringe sellers have an advantage in selling over M .

The rest of the paper proceeds as follows. In Section 1.1 we survey the related literature. We lay out the baseline model in Section 2 and analyze it in Section 3, where we compare the three modes that the platform can choose and the implications of banning the dual mode. In Section 4 we introduce product imitation and self-preferencing by the platform, and compare the four policy interventions mentioned above: banning the dual mode, banning imitation alone, banning self-preferencing alone, and banning both imitation and self-preferencing. Section 5 explores two main extensions of our framework: (i) exploring what happens when we vary how self-preferencing and imitation work in dual mode, including allowing for imitation to only happen a fraction of the time, allowing for imitation to be less than perfect, allowing imitation to be value-enhancing as would be the case if M can use its advantages to improve upon S 's innovation, and in case M can only steer a fraction of consumers; and (ii) comparing the

marketplace-seller dual mode analyzed here to the more traditional case in which the intermediary is a retailer (like a supermarket) that can offer its own in-house brands alongside products sourced from third-party suppliers. Finally, in Section 6 we conclude.

1.1 Related literature

A recent strand of literature has emerged that compares the platform business model with various alternative models: marketplace or reseller (Hagiu and Wright, 2015a), platform or vertically integrated firm (Hagiu and Wright, 2015b and 2019), agency or wholesale pricing (Johnson, 2017). In these papers, the key distinction between the business models is the delegation of control rights over key factors that are relevant for total demand, e.g. prices and marketing choices. This literature does not consider the possibility of the dual mode, in which a platform operates both as a marketplace and as a seller, competing with third-party sellers on its marketplace.

Somewhat closer is the literature that considers whether a platform should offer its own products or services.⁴ For example, Hagiu and Spulber (2013) consider a platform facing the chicken-and-egg coordination problem in user participation, showing that this problem can be mitigated by introducing first-party content alongside third-party content. Farrell and Katz (2000) and Jiang, Jerath, and Srinivasan (2011) analyze platform owners that face a tradeoff between extracting rents and motivating innovation by third-party complementors. Zhu and Liu (2018) empirically investigate this question, showing that Amazon is more likely to compete with its marketplace sellers in product categories that are more successful in terms of sales. These papers do not consider the surplus and welfare implications of the dual mode.

Our analysis of the dual mode with self-preferencing relates to the work by de Cornière and Taylor (2019), which considers a vertically integrated intermediary that biases its recommendations in favor of its subsidiary seller at the expense of third-party sellers.⁵ In their setup, divestiture (which would eliminate the dual role played by the intermediary) means both the intermediary and the seller coexist and operate independently. Among several results, they show that divestiture can increase consumer surplus under price competition (as opposed to quantity competition).⁶ Our analysis of self-preferencing differs in many respects, including the following: we model the platform’s recommendation as a decision made after all prices are set, we allow for multiple channels (direct vs. intermediated), we allow for imitation, and we endogenize the intermediary’s choice of business model.

Concurrent with this paper, Etro (2021) investigates a platform’s choice between (i) earning commissions from third-party sellers that compete among themselves; and (ii) entering with private label products or as a reseller of third-party products. Whenever the platform enters, it completely eliminates competition between the product it sells and third-party sellers by engaging in an extreme form of self-preferencing. The paper shows that the platform’s incentive to enter aligns with consumers’ interests when sellers are perfectly competitive, but there is generally insufficient entry by the platform when sellers have market power. In contrast, we model the price competition between the platform and the third-party sellers. We also consider different regulatory options, and endogenize the intermediary’s post-ban choice of business model.

Even more recently, other authors have started analyzing further implications of the dual mode. Madsen and Vellodi (2021) consider a dynamic model with demand uncertainty, in which the platform can optimally commit to a product introduction policy that is contingent on the realized demand state

⁴Hagiu et al. (2020) consider the opposite situation of a traditional firm hosting rivals to become a platform.

⁵See also the first section of Calvano and Polo (2021) for a comprehensive survey on the economic literature of biased intermediation by digital platforms.

⁶See also Gilbert (2020), Krämer and Zierke (2020), and Zenny (2020) on how vertical separation or banning self-preferencing can sometimes decrease surplus and efficiency in similar settings.

(which the platform observes from the marketplace data). They show that a regulation which bans the platform’s usage of marketplace data can either stifle or stimulate innovation, depending on the nature of innovation. Anderson and Bedre-Defolie (2021) consider the implications of the dual mode in a setting in which the platform is a price leader that competes against horizontally differentiated fringe sellers. Their work complements ours by highlighting a potential harm of the dual mode when there is a variety effect. They find that the platform sets a higher commission in the dual mode than in the pure marketplace mode, leading to fewer participating sellers and fewer product varieties, thus harming consumers. Their analysis utilizes the technique of aggregative games by assuming that consumers’ preferences follow the Gumbel distribution, which implies a constant markup enjoyed by sellers. This assumption shuts down the channel through which the dual mode reduces sellers’ market power, which is the main benefit of the dual mode in our model. Furthermore, our model also explores how the practices of product imitation and self-preferencing affect the implications of banning the dual mode.

At a high level, a platform that operates in dual mode can be viewed as a vertically integrated firm that uses the upstream input (the facilitation of transactions through its marketplace) to offer a downstream product (sell its own product through the marketplace) that competes with other downstream sellers. The literature on vertical foreclosure has studied how upstream market power leads to negative effects on downstream competition (e.g. Rey and Tirole, 2007). Our setting is different in several respects, including that the platform is not essential for fringe sellers and that all third-party sellers have the possibility to induce consumers to buy directly, possibilities that do not arise in traditional vertical foreclosure settings. Similarly, the literature on access pricing has explored how a vertically integrated incumbent distorts its competition with downstream firms through inefficient demand-sabotage (e.g. Mandy and Sappington, 2007; Brito et al., 2012). While demand-sabotage is analogous to self-preferencing under the dual mode in our analysis, one key conceptual difference is that self-preferencing does not directly reduce efficiency because consumers still purchase the product with the highest net value on the equilibrium path. Rather, self-preferencing generates an inefficiency only through its effect on the third-party’s incentive to innovate as opposed to any direct inefficiency associated with an actual act of sabotage.

2 Baseline model setup

Suppose each consumer wants to buy one unit of one product where there is a continuum (measure one) of consumers. Transactions can be performed directly or through a platform (or more generally, an intermediary) M . The direct channel corresponds to buying from the seller’s own website, or more generally through some alternative channel (including traditional stores in some cases). Consumers enjoy a convenience benefit $b > 0$ of using M to conduct transactions.

The product is supplied by at least two identical “fringe sellers”, and the products are each valued at v by consumers. In addition, there is a superior seller S which benefits from an innovation, such that its product is valued at $v + \Delta > v > 0$. Depending on the mode of operation, M may be able to operate as a seller and sell its “own offering” to consumers which is valued at $v + \sigma$, where σ can be positive or negative.⁷ Thus, consumers are homogenous with respect to their willingness to pay.⁸

Consumers always have an outside option of not buying anything, which gives them a surplus of v_o . We assume that v_o is distributed according to a log-concave and continuously differentiable cumulative

⁷This covers two possibilities which we will not distinguish: M ’s offering may be an existing product sourced from competitive manufacturers or wholesalers, or M ’s offering may be produced and sold under its own brand (like Amazon Basics or Apple’s own apps).

⁸In Section A of the Online Appendix, we consider an extension where consumers have heterogenous values for S ’s product. As long as there is not too much consumer heterogeneity, we continue to have similar equilibrium characterizations as in Section 3. The main implications of banning the dual mode in Proposition 4 still apply, except that the magnitude of any decrease in CS and W is smaller than in the case with no heterogeneity.

distribution function G on the support $[0, \bar{v}_o]$, where $\bar{v}_o \leq \infty$, so we allow distributions with bounded and unbounded support. Denote the corresponding density function as g , so that the log-concavity assumption implies $\frac{d}{dx} \frac{G(x)}{g(x)} \geq 0$. All marginal costs are normalized to zero.

□ **Innovation.** S can choose the level of its innovation $\Delta \geq \Delta^l$ by incurring a fixed cost $K(\Delta)$, where $\Delta^l > 0$ is the default innovation level of S 's product. The cost function $K(\cdot)$ is increasing and convex, and its derivative is denoted as $K'(\cdot)$. We assume $K(\Delta^l) = 0$ and $K'(\Delta^l) \leq G(v)$ (which is a more general version of the usual boundary condition that $K'(\Delta^l) = 0$). Denote $\bar{\Delta} > \Delta^l$ as the solution to the following first-order condition:

$$K'(\bar{\Delta}) = G(v + \sigma + b). \quad (1)$$

Throughout the paper, $\bar{\Delta}$ is the highest possible innovation level that will arise in equilibrium (in all settings throughout the paper). In the baseline version of the model we assume that M cannot copy S 's innovation.

□ **Product discovery on the marketplace.** All consumers are initially aware of all the offerings in the market except S 's product. Consumers can discover S through M , provided that S is available on M 's marketplace. In practice, many consumers rely on marketplaces like Amazon and Apple's App Store to find out about the existence of new products. In the baseline version of the model we assume that M does not engage in self-preferencing, so that all consumers become aware of S 's existence as long as S participates on the marketplace. Once a consumer becomes aware of S 's existence through M , she also becomes aware of S 's direct channel and its associated price. She is free to choose which channel to buy S 's product from (and still enjoy transaction benefit b if she buys it through M 's marketplace).

□ **Platform commission.** Whenever M 's mode includes a marketplace, it charges a commission $\tau \geq 0$ to third-party sellers for each transaction facilitated.⁹ Third-party sellers (including S) can choose whether to participate on M 's marketplace, and whenever they do, can price discriminate between consumers that come to it through the marketplace and consumers that come to it through the direct channel. We posit that third-party sellers participate on M 's marketplace whenever they are indifferent.

Given there are always two or more identical fringe sellers competing in the direct channel and on the marketplace, following the standard Bertrand logic, we take as given that fringe sellers always price at marginal cost, i.e. zero if selling directly and τ if selling on a marketplace, regardless of how S and M price. Thus, when we characterize equilibria, we take these fringe seller prices as given.¹⁰

□ **Timing:**

1. M chooses its mode of operation and sets τ (if in marketplace or dual mode);
2. S chooses whether to participate on the platform (if applicable) and the level of its innovation Δ ;
3. All sellers, including S and M , set prices simultaneously;
4. Consumers make their purchase decisions (after observing the existence of S in case it participates on the platform).

We solve for Subgame Perfect Nash Equilibria. Whenever there are multiple equilibria in any subgame that are payoff-ranked by M , we select the one preferred by M . Then, whenever there are multiple

⁹We assume the transaction fee is fixed per transaction rather than being a percentage of the value of the transaction for simplicity. With percentage fees, which are more commonly observed, S 's optimal innovation would continuously depend on the fee level (in both the marketplace and the dual modes). The resulting analysis is no longer tractable. Nevertheless, in Section B of the Online Appendix we fix the level of S 's innovation, and show that the key insights of the baseline model remain valid when the platform charges percentage fees.

¹⁰Note this implies that throughout the paper we are ruling out equilibria supported by fringe suppliers pricing below cost, i.e. playing weakly dominated strategies.

equilibria in any subgame that are payoff-equivalent for M , but payoff-ranked by S , we select the one preferred by S . Throughout, we rule out equilibria in any subgames which rely on firms playing weakly dominated strategies.

It is useful to discuss the interpretation of the parameter σ in our model. While we will refer to $\sigma > 0$ as increasing the value of M 's offering, this is just for expositional simplicity, and we could have equivalently defined it as the cost advantage that M enjoys when selling the same product as third-party sellers. In practice, $\sigma > 0$ could reflect the superiority of the platform's own offering in terms of reputation, customer service and logistics, lower input prices due to economies of scale, or in the case of apps, better integration with the core platform services. Meanwhile, $\sigma < 0$ could reflect the existence of niche product knowledge or technical know-how that third-party sellers may have over the platform. Note σ plays a different role from b in the analysis. The latter ensures that M can sustain positive transaction fees in its marketplace given the ability of consumers to buy outside. The former allows us to have a tradeoff between the seller and marketplace modes, since it captures the advantage or disadvantage that M has in selling relative to third-parties.

We assume $\sigma < \bar{\Delta}$, i.e., S 's product can be more valuable than M 's, as otherwise in dual mode there can be no equilibrium where S makes any sales via M , and the dual mode simply reduces to the seller mode. We also assume $b + \sigma > 0$ as otherwise M 's own offering is worse than the fringe's product in the direct channel, which implies that M 's seller mode is not viable and that the dual mode reduces to the marketplace mode. Finally, we assume

$$\max\{b, b + \sigma, \bar{\Delta}\} < \frac{G(v)}{g(v)} \quad (2)$$

to focus on the interesting case where the equilibrium prices are always determined by the competitive constraint. To ensure equilibria are well defined, we assume that consumers break ties in favor of S 's product or M 's offering whenever they are indifferent between these and the fringe sellers' products.

3 Baseline analysis

In this section we characterize the equilibria arising after M 's choice of each of the three possible modes: marketplace mode, seller mode, or dual mode, as well as the consequences of banning the dual mode. In the first period, M chooses one of these modes, and this becomes common knowledge.¹¹

3.1 Marketplace mode

Suppose M chooses the marketplace mode. Consider S 's pricing decision in stage 3 after it has joined the marketplace and chosen innovation level Δ . Clearly, S 's price on the marketplace ("inside" price) p_i and direct price ("outside" price) p_o are bounded above by the fringe sellers' prices on M after accounting for the superior quality of S 's product, i.e. $p_i \leq \tau + \Delta$ and $p_o \leq \tau + \Delta - b$. Notice that S can influence consumers' transaction channel by adjusting the relative level of its inside and outside prices. Specifically, S adopts one of the following strategies:

- Set $p_o > p_i - b$ (so that any consumer buying from S does so through M) and p_i that solves

$$\max_{p_i \leq \tau + \Delta} (p_i - \tau)G(v + b + \Delta - p_i). \quad (3)$$

¹¹Implicit in this timing assumption is that if it chooses the marketplace mode, M is committed not to compete with third-party sellers after they participate. In practice, this commitment could be established by building a reputation for remaining a pure marketplace, as eBay has done since its inception.

- Set $p_i > p_o + b$ (so that any consumer buying from S does so directly) and p_o that solves $\max_{p_o \leq \tau + \Delta - b} p_o G(v + \Delta - p_o)$.

It is easy to check that S optimally chooses the first strategy if $\tau \leq b$ and the second strategy if $\tau > b$. This reflects the “showrooming constraint” (to use Wang and Wright (2020)’s terminology), whereby the platform can only set its transaction fee up to the transaction convenient benefit b it offers. Any $\tau > b$ would result in S joining M (to make all consumers aware of its existence) and then setting prices to induce all consumers to purchase directly, so that M earns zero profit.

Thus, in equilibrium M sets $\tau \leq b$. After taking into account S ’s endogenous innovation, we have the following equilibrium characterization:

Proposition 1 (*Marketplace mode equilibrium*) M sets $\tau^{mkt} = b$, while S participates, sets $\Delta = \Delta^{mkt}$ as defined by

$$G(v) = K'(\Delta^{mkt}), \quad (4)$$

and sells exclusively through the marketplace at price $p_i^* = b + \Delta^{mkt}$.

Intuitively, (2) implies that: (i) the competitive constraint binds in (3), so $p_i^* = \tau^{mkt} + \Delta$ in equilibrium; (ii) the showrooming constraint on M ’s fee binds, so M optimally chooses $\tau^{mkt} = b$. Thus, the equilibrium profits of M and S are $\Pi^{mkt} = bG(v)$ and $\pi^{mkt} = \Delta^{mkt}G(v) - K(\Delta^{mkt})$.

3.2 Seller mode

Suppose M chooses the seller mode and sets a price p_m for the product it sells. Given that S is unavailable on M and so it is never shown to any consumer, S ’s price is irrelevant. It sets any $p_o^* \geq 0$ and chooses the lowest possible innovation level $\Delta^{sell} = \Delta^l$ in equilibrium. Then, M ’s price is only constrained by competition with fringe sellers selling through their direct channels. It solves

$$\max_{p_m \leq b + \sigma} p_m G(v + b + \sigma - p_m).$$

Assumption (2) implies that the competitive constraint binds, so $p_m^* = b + \sigma \geq 0$. The equilibrium profits of M and S are $\Pi^{sell} = (b + \sigma)G(v)$ and $\pi^{sell} = 0$.

Proposition 2 (*Seller mode equilibrium*) M sets $p_m^* = b + \sigma$ and sells to all consumers, while S sells to no one.

3.3 Dual mode

Suppose M chooses the dual mode, and consider the pricing subgame in stage 3. Let us focus on $\tau \in (\max\{-\sigma, 0\}, b]$, which turns out to be part of the overall equilibrium when M sets τ . Then, there are two possible equilibria in the pricing subgame:

- (*Semi-seller mode equilibrium*) If $\sigma \geq \Delta$, all consumers buy from M . The equilibrium prices are $p_i^* = \tau$, $p_o^* \geq p_i^* - b$, and $p_m^* = \tau + \sigma - \Delta$.
- (*Price squeeze equilibrium*) If $\Delta > \sigma$, all consumers buy from S through the marketplace. Any price profile satisfying $p_i^* = p_m^* + \Delta - \sigma$, $p_o^* \geq p_i^* - b$, and

$$p_m^* \in [\max\{\tau - \Delta + \sigma, 0\}, \tau + \min\{\sigma, 0\}] \quad (5)$$

is an equilibrium.

The semi-seller mode equilibrium reflects the standard asymmetric-good Bertrand competition. Given $\sigma \geq \Delta$, M has the superior product and so it sells to all consumers. Meanwhile, S sets its price at its effective marginal cost τ and makes zero profit.

The more novel case is the price squeeze equilibrium. A few remarks are in order on the construction of this equilibrium. First, M 's equilibrium inside price must satisfy $p_m^* \leq \tau + \min\{\sigma, 0\}$. If $p_m^* > \tau$, M could earn a higher margin than τ by undercutting S and selling its own offering. If $p_m^* > \tau + \sigma$ consumers prefer the fringe product on the marketplace over M 's offering. Second, $p_m^* \leq \tau + \min\{\sigma, 0\}$ implies M 's price imposes a stronger constraint on S 's price than do the fringe sellers. This is the sense in which M exerts a “price squeeze” on S . As such, S 's pricing problem is

$$\max_{p_i \leq p_m^* + \Delta - \sigma} (p_i - \tau)G(v + b + \Delta - p_i),$$

and (2) implies that S 's pricing constraint must bind, so it sets $p_i^* = p_m^* + \Delta - \sigma$. Finally, given that the effective marginal costs of S and M are τ and 0 respectively, the equilibrium must satisfy $p_i^* \geq \tau$ and $p_m^* \geq 0$, which gives the lower bound in (5).¹² Even though M 's marginal cost is zero, it does not necessarily want to undercut S because $p_m^* \leq \tau$ implies that undercutting leads to a smaller margin than it can get from its commission τ .

Notice that there are multiple price squeeze equilibria. For each p_m^* in (5), the equilibrium profits of M is

$$\Pi = \tau G(v + \sigma + b - p_m^*). \quad (6)$$

Given our equilibrium selection rule, we select the lowest price $p_m^* = \max\{\tau - \Delta + \sigma, 0\}$, which maximizes M 's profit in (6).¹³

□ **Stage-2 innovation decision.** Taking into account both of the possible equilibria in the pricing subgame, we obtain the following lemma on S 's innovation decision.

Lemma 1 Denote $\bar{\tau} \in (\Delta^l - \sigma, \bar{\Delta} - \sigma)$ as the unique solution of

$$(\bar{\Delta} - \sigma - \bar{\tau})G(v + b + \sigma) - K(\bar{\Delta}) = 0. \quad (7)$$

In stage 2, S sets $\Delta = \bar{\Delta}$ if $\tau \leq \bar{\tau}$, and sets $\Delta = \Delta^l$ if $\tau > \bar{\tau}$.

The term $\bar{\tau}$ in Lemma 1 plays an important role in equilibrium, and it is the highest commission such that S is still willing to innovate (i.e., choosing $\Delta = \bar{\Delta} > \Delta^l$). We will refer to the constraint that $\tau \leq \bar{\tau}$ as the “*innovation constraint*”. For $\tau > \bar{\tau}$, the commission is too high such that S simply sets the lowest possible innovation level Δ^l to save on the innovation fixed cost and earns a zero margin. For $\tau \leq \bar{\tau}$, the commission is low relative to the innovation cost. S sets a high innovation level that generates a strictly positive margin $p_i^* - \tau = \bar{\Delta} - \sigma - \tau > 0$ in the resulting price squeeze equilibrium. Reflecting this logic, notice from (7) that $\bar{\tau}$ is higher if $\bar{\Delta}$ is large relative to $K(\bar{\Delta})$, i.e., when innovation is cost-efficient.

□ **Stage-1 commission decision.** M 's choice of commission reflects an interaction between the “showrooming constraint” ($\tau \leq b$) and the “innovation constraint” ($\tau \leq \bar{\tau}$), as described in the result below:

¹²Strictly speaking, there can be other equilibria involving $p_m < 0$. However, for any given τ , these equilibria involve M playing weakly dominated strategies. We rule out such equilibria because they involve M setting a price that it would prefer to change if some consumers actually purchased from it (i.e. off the equilibrium path).

¹³In Section C of the Online Appendix, we show that the prices implied by our equilibrium selection rule are also the unique equilibrium outcome if instead we had assumed that M was the price leader and could commit to its price before S .

Proposition 3 (*Dual mode equilibrium*)

- If $b \leq \bar{\tau}$ or

$$\bar{\tau}G(v + \sigma + b) \geq (b + \max\{\sigma - \Delta^l, 0\})G(v + \Delta^l), \quad (8)$$

then M sets $\tau^{dual} = \min\{b, \bar{\tau}\}$, while S participates, sets $\Delta^{dual} = \bar{\Delta}$ and sells to all consumers through the marketplace at price $p_i^* = \bar{\Delta} - \sigma$.

- If $b > \bar{\tau}$ and (8) does not hold, then M sets $\tau^{dual} = b$, while S participates and sets $\Delta^{dual} = \Delta^l$. If $\Delta^l > \sigma$, then S sells to all consumers exclusively through the marketplace at price $p_i^* = \tau^{dual}$. If $\Delta^l \leq \sigma$, then M sells to all consumers at price $p_m^* = \tau^{dual} + \sigma - \Delta^l$.

Intuitively, if $b \leq \bar{\tau}$, then S chooses a high innovation level for all $\tau \leq b$ (Lemma 1), so that the showrooming constraint is the only binding constraint on M 's fee ($\tau^{dual} = b$), with M 's equilibrium profit being $\Pi^{dual} = bG(v + \sigma + b)$.

If $b > \bar{\tau}$, then M faces a trade-off that is captured by (8). Satisfying the innovation constraint at $\tau = \bar{\tau} < b$ reduces M 's margin but encourages S 's innovation, which intensifies the on-platform competition and results in more transactions. Violating the innovation constraint at a high fee $\tau = b$ increases M 's margin but discourages S 's innovation, which leads to fewer transactions. In this case, Π^{dual} equals the side which is larger in the inequality in (8).

Notice that in equilibrium, competition with M on the marketplace effectively imposes a “price squeeze” on S 's inside price. By setting a low p_m , M induces a lower equilibrium inside price by S , which leads to more transactions through the marketplace. In the first case where $\Delta^{dual} = \bar{\Delta}$, we have $p_i^* = \bar{\Delta} - \sigma \geq \tau^{dual}$, meaning that the price squeeze is only partial, in the sense that S still earns a positive margin. In this case, the equilibrium innovation level $\bar{\Delta}$ is high enough such that if M wants to cap S 's inside price further, then M would need to set $p_m < 0$. Such a price below marginal cost is a weakly dominated strategy in the pricing subgame, which is why we ruled it out, and so the tightest possible price cap M can impose on S (achieved when $p_m = 0$) is $p_i^* = \bar{\Delta} - \sigma > \tau^{dual}$. The second case where $\Delta^{dual} = \Delta^l$ reflects that the price squeeze is sufficiently strong or innovation is sufficiently costly, so that S chooses the lowest innovation level Δ^l . The price squeeze is complete in the sense that M caps S 's inside price to $p_i^* = \tau^{dual}$ while still setting a non-negative p_m .

3.4 Choice of mode

We are now ready to compare M 's profits across all three modes.

Corollary 1 (*Platform profit*)

- $\Pi^{mkt} \geq \Pi^{sell}$ if and only if $\sigma \leq 0$, with equality holding only when $\sigma = 0$;
- $\Pi^{dual} > \Pi^{mkt}$;
- There exists a unique threshold $\underline{\sigma} > 0$ such that $\Pi^{dual} > \Pi^{sell}$ if and only if $\sigma < \underline{\sigma}$.

Comparing the two pure modes, M prefers the marketplace mode if $\sigma < 0$ and the seller mode if $\sigma > 0$. In the marketplace mode, M hosts S 's innovative product, which allows M 's channel (as a whole) to compete more favorably against fringe sellers in their direct channel. However, M can only partially extract the resulting sales revenue through its commission due to the existence of the showrooming constraint. In contrast, in the seller mode, M fully extracts any sales revenue from its channel, but M 's

inferior product means the competition with fringe sellers is less favorable than in the marketplace mode when $\sigma \leq 0$.

Following the price squeeze logic in the dual mode, $\Pi^{dual} > \Pi^{mkt}$ because S 's lower price leads to more transactions on the marketplace. Due to this extensive margin, the dual mode is strictly more profitable even though the equilibrium commission is the same across both modes.

Finally, the comparison between Π^{dual} and Π^{sell} (when $\sigma > 0$) reflects two opposing forces when M opens up its channel for S to make sales in dual mode. First, by having S 's superior product on the marketplace and squeezing S 's price, M can generate more transactions. Second, given that M does not steer in this baseline setup, by choosing dual mode, M essentially commits to let all consumers become aware of S 's product. This means that M can no longer fully exploit its own competitive advantage σ due to competition with S . The dual mode is less profitable than the seller mode if σ is sufficiently large. This is consistent with the result in Hagiou et al. (2020), in which a platform only wants to host a rival when the rival's product is sufficiently superior to its own.

3.5 Banning dual mode in the baseline model

A policy that bans the dual mode can result in two possible market structures, depending on whether M chooses to operate in the marketplace mode or the seller mode in period zero (Corollary 1). We examine the implications on profits (Π and π), consumer surplus (CS), innovation (Δ), and welfare (W).

Proposition 4 *Banning the dual mode in the baseline model has the following effects:*

	<i>M's choice of mode</i>	Π	π	CS	Δ	W
<i>if $\sigma \geq \underline{\sigma}$</i>	<i>Seller</i>
<i>if $\sigma \in (0, \underline{\sigma})$</i>	<i>Seller</i>	↓	↓ <i>if $b < \bar{\tau}$;</i> . <i>if $b \geq \bar{\tau}$</i>	↓	↓ <i>if $b \leq \bar{\tau}$ or</i> <i>(8) holds;</i> . <i>otherwise</i>	↓
<i>if $\sigma \leq 0$</i>	<i>Marketplace</i>	↓	↑	↓	↓ <i>if $b \leq \bar{\tau}$ or</i> <i>(8) holds;</i> ↑ <i>otherwise</i>	↓ <i>if Δ decreases or</i> <i>(9) does not hold;</i> ↑ <i>otherwise</i>

“.” = not changing; “↑” = increasing; “↓” = decreasing.

For $\sigma \geq \underline{\sigma}$, banning the dual mode has no effect as M always operates in the seller mode.

For $\sigma \in (0, \underline{\sigma})$, M switches from the dual mode to the seller mode after the ban. Consumer surplus decreases because price competition is weak when M operates in the seller mode, given that consumers are unaware of S 's existence. The corresponding decrease in welfare reflects two sources of inefficiencies in the seller mode: (i) there are fewer transactions; and (ii) consumers can no longer combine S 's superior product with M 's convenience benefit.

For $\sigma \leq 0$, M switches from dual mode to the marketplace mode after the ban. Consumer surplus decreases due to the weaker on-platform competition in the marketplace mode. Innovation is lower in the marketplace mode if $b \leq \bar{\tau}$ or (8) holds, so that $\Delta^{dual} = \bar{\Delta} > \Delta^{mkt}$. This is because the partial price squeeze in dual mode leads to a greater transaction volume, which increases the marginal benefit of S 's innovation and enhances S 's innovation incentive (compare (4) and (1)). However, innovation is higher in the marketplace mode if $b > \bar{\tau}$ and (8) fails, whereby $\Delta^{dual} = \Delta^l$. In this case, the price squeeze in dual mode is too strong relative to the cost-efficiency of S 's innovation, so that S has no innovation incentive in dual mode.

Finally, the change in welfare after the switch to the marketplace mode reflects two opposing forces: (i) there are fewer transactions in marketplace mode due to the higher prices that result from S facing

less competition; and (ii) S has a higher innovation incentive in the marketplace mode if $\Delta^{dual} = \Delta^l$. Thus, welfare is higher in the marketplace mode if and only if both $\Delta^{dual} = \Delta^l$ and

$$(b + \Delta^{mkt})G(v) - K(\Delta^{mkt}) > bG(v + \Delta^l) + \int_0^{v+\Delta^l} [v + \Delta^l - \max\{v, v_o\}]dG(v_o) \quad (9)$$

hold simultaneously.

To summarize, Proposition 4 shows that in this baseline setting, a ban on dual mode, whenever it is relevant, always results in lower consumer surplus. The ban also decreases total welfare when $\sigma \in (0, \underline{\sigma})$, i.e. when M has an advantage in selling its product over fringe sellers. Even when M has a disadvantage in selling its product ($\sigma \leq 0$), the ban does not necessarily increase total welfare. For that to be the case, we also require that the ban increases innovation. From Proposition 4, it can be shown that a necessary condition for such an increase is $\Delta^l - b < \sigma \leq 0$, i.e., the convenience benefit b is strictly greater than S 's advantage Δ^l .

It is worth emphasizing that for large platforms like Amazon and Apple, the case $\sigma \leq 0$ is likely less empirically relevant than the case $\sigma > 0$. Amazon and Apple benefit from significant economies of scale and scope, which should allow them to produce more efficiently than fringe sellers (either better products for the same cost, which is the way we have interpreted $\sigma > 0$, or equivalently, the same products at lower costs). However, the platforms' products may not be as good as the versions produced by sufficiently innovative sellers.

4 Product imitation and self-preferencing

In this section we explore the two practices that certain platforms that operate in dual mode (such as Amazon and Apple) have been alleged to use, and that have drawn scrutiny by policymakers: imitation of third-party products and self-preferencing (steering consumers to the products sold by the platform). We enrich the baseline model by introducing both of these practices below.

□ **Product imitation.** Whenever S is available on the platform and M operates in dual mode, M (before setting its price) can choose to imitate S 's superior product, thereby also offering consumers a product with the same surplus as S 's product, i.e. $v + \Delta$. Note M could still choose to offer its original product as well, which it would do if and only if $\sigma \geq \Delta$. Thus, here we focus on the strongest form of imitation in dual mode—it occurs immediately and works perfectly. If instead S does not participate, then M cannot imitate S 's product. This captures the policy concern that by hosting third-party sellers, M obtains some kind of proprietary data from them which allows it to copy the relevant product features from that seller. Implicit in this timing assumption is that M cannot commit to not imitate S 's product. To simplify the exposition, we will assume that M breaks ties in favor of not imitating whenever it is indifferent, reflecting that imitation may be costly.

□ **Steering.** To model the possibility of M engaging in self-preferencing, we assume that all consumers rely on M 's recommendation to discover S 's novel product, so that M can steer consumers by determining whether or not they are aware of S 's existence (e.g., through its recommendation algorithm). In other words, unlike in the baseline model, S 's participation on M in the marketplace and dual modes is no longer sufficient for consumers to learn about its existence. Specifically, after all prices are set, M makes a binary choice of whether to show S 's product to consumers or not. Consumers remain aware of all other products, including any product that M is selling.

□ **Timing.** The timing of this enriched model is exactly the same as in the baseline model, except the details in Stages 2 and 4 are modified as follows to handle the possibility of imitation and steering:

- Stage 2: a) S chooses whether or not to participate on the platform (if applicable) and the level of

its innovation Δ ; b) if M is operating in dual mode and S participates, M chooses whether or not to imitate S 's product.

- Stage 4: a) M chooses whether or not to show S (if S participates on the platform); b) After observing the existence of S if M shows it, all consumers make their purchase decisions.

Before proceeding, we note that this enriched model is set up in a way that maximizes the potential harm (to innovation, consumer surplus and welfare) that occurs in dual mode. Indeed, in dual mode, the platform can steer perfectly, engage in perfect product imitation, and is unable to commit not to imitate. As such, one can interpret the results in this section as identifying the worst-case scenario associated with self-preferencing and imitation.

In Section 5, we consider imperfect steering, imperfect or value-adding product imitation, and the possibility that the platform can commit not to imitate. Each of these extensions complicates the analysis, but brings the results back closer to the baseline results in Section 3.5. This suggests that a more realistic case lies somewhere between the baseline results and those presented in this section.

4.1 Pure modes

□ **Marketplace mode.** We first derive M 's stage-4 recommendation decision (whether to show S) after prices are set. Since M 's commission τ is the same for every seller, M shows S to consumers whenever doing so results in transactions on the marketplace. If M shows S , consumers buy from S through the marketplace (provided that they buy anything at all) if and only if (i) consumers do not switch to purchasing directly after learning of S 's existence; (ii) consumers prefer buying from S through the marketplace rather than buying from fringe sellers (through either channel). Formally:

$$\Delta + b - p_i \geq \max \{ \Delta - p_o, b - \tau, 0 \}. \quad (10)$$

If M does not show S , consumers buy the fringe product through the marketplace if and only if

$$b - \tau \geq 0. \quad (11)$$

The optimal recommendation rule is straightforward when exactly one of the conditions (10) and (11) holds. However, when (10) and (11) hold simultaneously, M is indifferent between recommending S or not because consumers always buy from the marketplace, so M always obtains τ . Likewise, when (10) and (11) fail simultaneously, M always obtains zero profit regardless of its recommendation. In such cases (which also arise below in dual mode), we assume that M breaks the tie in favor of showing S . This tie-breaking rule is consistent with our equilibrium selection rule, and it can also be interpreted as reflecting that any information manipulation (not showing S) may involve a small but non-zero cost for M .¹⁴

In the appendix, we show that this recommendation rule implies that the equilibrium outcome in marketplace mode is exactly the same as in the baseline model, i.e., Proposition 1 applies.

□ **Seller mode.** Neither of the new modelling ingredients (imitation and steering) has any effect on the seller mode, so the equilibrium in Proposition 2 applies.

¹⁴This cost could capture the risks that arise from any such steering, including damage to M 's reputation and regulatory intervention. Alternatively, S could offer a small additional commission to M in order to be shown given that the tie-breaking rule affects S 's profit but does not affect M 's profit.

4.2 Dual mode with product imitation and self-preferencing

To keep the exposition concise, in what follows we focus on presenting the main qualitative features of the analysis and we relegate more formal details to the appendix. Given the possibility of product imitation, we assume

$$b + \bar{\Delta} < \frac{G(v)}{g(v)}. \quad (12)$$

This serves the same purpose as (2) did in the baseline, i.e. it ensures that the competitive constraints are always binding on equilibrium prices.

□ **Recommendation.** Given that M is selling (either its own product or its imitation of S 's product), its recommendation decision is determined by comparing its own margin with the commission, after adjusting for the probabilities of consumers purchasing each of the products (given that consumers have heterogenous outside options). In stage 4, M optimally chooses not to show S if at least one of the following conditions holds:

- M 's expected margin from selling (whether its own product or its imitation of S 's product) is higher than the expected commission it could collect from S and consumers prefer to buy M 's offering over buying from the fringe sellers (either directly or on M).
- Consumers do not buy from S through the marketplace when S is shown, i.e., $\Delta + b - p_i < \max\{\Delta - p_o, b - \tau, 0\}$.

If neither of these conditions holds, then M 's expected margin is lower than the expected commission, and showing S results in transactions on the marketplace. In this case, M will show S 's product to consumers.

□ **Pricing subgame without imitation.** For the stage-3 pricing subgame, we first consider the case where M has chosen not to imitate S 's product. This allows us to understand how M 's ability to steer affects the pricing subgame relative to the baseline in Section 3.3. Broadly speaking, there are two types of relevant equilibria in the subgame, depending on the value of τ (the complete equilibrium strategies are stated in the proof of Proposition 5).

- *Exploitative equilibrium (without imitation).* This parallels the semi-seller mode equilibrium in the baseline dual mode. The key difference is that M can choose not to show S 's product in order to shield M 's own offering from competing with S . This allows M to fully extract the value of its original product σ , as well as $\min\{b, \tau\}$ (due to competition with fringe suppliers). Thus, M can set the “exploitative price” $p_m^* = \min\{\tau, b\} + \sigma$ and sell to all consumers, earning

$$\Pi_{no-imi}^{exploit} \equiv (\min\{\tau, b\} + \sigma)G(v + b - \min\{\tau, b\}). \quad (13)$$

- *Price squeeze equilibrium (without imitation).* The equilibrium construction for this case is similar to the baseline model, with M earning (6). However, here S may still sell through the marketplace in the price squeeze equilibrium even when $\tau > b$. This is because steering implies that S faces the threat of not being shown whenever it attempts to attract consumers to the direct channel.

□ **Pricing subgame with imitation.** Clearly, M has no incentive to imitate if $\sigma \geq \Delta$. So this pricing subgame is relevant only if $\Delta > \sigma$. In this case, after imitation, M 's product has the same value as S 's. Relative to the exploitative equilibrium without imitation, M can now set a higher exploitative price $p_m^* = \min\{\tau, b\} + \Delta$ and sell to all consumers, earning

$$\tilde{\Pi}_{imi}^{exploit} = (\min\{\tau, b\} + \Delta)G(v + b - \min\{\tau, b\}).$$

Meanwhile, the construction of the price squeeze equilibrium is unaffected, except that product imitation allows M to fully squeeze S 's price such that $p_m^* = p_i^* = \tau$ in this equilibrium (whenever it exists), regardless of τ .

Solving for the innovation, imitation, and commission decisions, the following proposition characterizes the overall equilibrium.

Proposition 5 (*Dual mode equilibrium with self-preferencing and imitation*). *If $\sigma \geq \Delta^l$, in equilibrium M sets $\tau^{dual} = b + \Delta^l$, then S participates and sets $\Delta^{dual} = \Delta^l$, then M chooses not to imitate and sells to all consumers at price $p_m^* = b + \sigma$. If $\sigma < \Delta^l$, both configurations below are equilibria with S participating and setting $\Delta^{dual} = \Delta^l$, and they are outcome-equivalent in terms of profits, consumer surplus, and welfare.*

- M sets $\tau^{dual} = b + \Delta^l$ and does not imitate. S sells to all consumers through the marketplace at price $p_i^* = \tau$.
- M sets $\tau^{dual} \neq b + \Delta$ and imitates. M sells to all consumers at price $p_m^* = b + \Delta^l$.

In equilibrium, S always chooses the lowest possible innovation Δ^l . This result is intuitive whenever its product is imitated by M on the equilibrium path. If instead M does not imitate S , then it must be the case that (i) M 's original product is superior ($\sigma \geq \Delta$); or (ii) τ is sufficiently high such that M can fully extract S 's margin without resorting to imitation. Both possibilities imply S 's profit is always $-K(\Delta)$ and so it always chooses Δ^l .

In Proposition 5, the case $\sigma \geq \Delta^l$ results in the exploitative equilibrium (without imitation) with M selling to all consumers and earning

$$\Pi^{dual} = (b + \sigma)G(v).$$

Self-preferencing means M faces no competitive pressure from S , so everything is as if M were a pure seller.

The case $\Delta^l > \sigma$ results in M fully extracting the innovation surplus of S 's innovative product through a combination of product imitation and self-preferencing, so M earns

$$\Pi^{dual} = (b + \Delta^l)G(v).$$

There are two ways in which M can achieve this profit. First, M can set $\tau = b + \Delta$ to induce the price squeeze equilibrium (without imitation), in which case S sells to all consumers but M fully extracts S 's profit through a high commission. This is possible because off-equilibrium path steering relaxes the showrooming constraint on M 's commission. Even though there is no imitation on the equilibrium path in this case, imitation off the equilibrium path still eliminates S 's innovation incentive. Second, M can set any $\tau \geq b$ such that $\tau \neq b + \Delta$ to induce the exploitative equilibrium (with imitation). M sells the imitation product to all consumers at a high price. This is possible because self-preferencing shields M 's product from competition with S .

The existence of the two outcome-equivalent equilibria when $\Delta^l > \sigma$ is a consequence of the assumption that M 's imitation product has the same value as S 's product. In Section 5, we discuss an extension of our model in which M 's imitation product is either imperfect (offers a value below S 's) or value-adding (offers a value above S 's). If imitation is imperfect, M strictly prefers the first approach of inducing the price squeeze equilibrium without imitation. If imitation is value-adding, M strictly prefers the second approach of inducing the exploitative equilibrium (with imitation). Our focus on perfect imitation is thus a limit case that lies between these two alternatives. Despite the discrete switch in M 's equilibrium

strategy when we go from one alternative to the other, we note that due to the equivalence result under perfect imitation, the market outcome (profits, consumer surplus, and welfare) is continuous across the cases of imperfect imitation and value-adding imitation.

In sum, Proposition 5 highlights two potential downsides of the dual mode when perfect self-preferencing and perfect imitation are possible: (i) the lack of on-platform competition; and (ii) the elimination of S 's innovation incentive. In the next section, we consider two distinct policy approaches to address these potential downsides.

4.3 Banning dual mode with self-preferencing and imitation

Consider first the structural approach of banning the dual mode.

Proposition 6 *Banning the dual mode (with self-preferencing and imitation) has the following effects:*

	M 's choice of mode	Π	π	CS	Δ	W
if $\sigma \geq \Delta^l$	<i>Seller</i>
if $0 < \sigma < \Delta^l$	<i>Seller</i>	↓	.	.	.	↓
if $\sigma \leq 0$	<i>Marketplace</i>	↓	↑	.	↑	↑

“.” = not changing; “↑” = increasing; “↓” = decreasing.

The results in Proposition 6 contain two major differences relative to Proposition 4 in the baseline setup. First, when $\sigma > 0$, banning the dual mode (which results in a switch to the seller mode) does not affect consumer surplus, as opposed to decreasing it in the baseline setup. This is because steering by M implies a high final consumer price in dual mode, either by enabling a high commission or by weakening on-platform competition. As such, the switch to the seller mode does not lead to any further increase in the price level. However, banning the dual mode still lowers welfare provided S 's product is superior to M 's and M still has an advantage in selling over fringe sellers (i.e. $0 < \sigma < \Delta^l$).

Second, in the empirically less likely case that M actually has a disadvantage in selling compared to fringe sellers (i.e. $\sigma \leq 0$), banning the dual mode does not affect consumer surplus (as opposed to decreasing it in the baseline setup) and always increases innovation and welfare (as opposed to being ambiguous). Indeed, we know that M 's ability to perfectly steer consumers means that the final price is the same in the marketplace mode and in the dual mode, explaining why consumer surplus does not change. At the same time, the possibility of imitation by M (on and off the equilibrium path) in dual mode eliminates S 's incentive to innovate so it always chooses Δ^l in dual mode.

4.4 Behavioral remedies

Instead of banning the dual mode, an alternative approach is to impose behavioral remedies such as (i) banning self-preferencing; (ii) banning imitation; or (iii) banning both self-preferencing and imitation. We consider the implications of each of these remedies.

□ **Ban on self-preferencing only.** A ban on self-preferencing requires M to always show S 's product when S is listed on its marketplace, so that all consumers are fully aware of all offers available in the market whenever S participates. There are several direct implications for the dual mode: (i) the exploitative equilibrium no longer exists because M can no longer avoid competing with S ; and (ii) the showrooming constraint on τ is restored. However, product imitation is still allowed. In the proof of the next proposition, we show that the post-intervention equilibrium of the dual mode is characterized by:

- M sets $\tau = b$, S participates and sets $\Delta = \Delta^l$, and M imitates. S sells to all consumers through the marketplace at $p_i^* = b$. Profits are $\Pi = bG(v + \Delta^l)$ and $\pi = 0$.¹⁵

Note that M 's post-intervention profit in dual mode is sometimes lower than the seller mode profit, especially when σ is high. In particular, there exists a cutoff $\underline{\sigma}^{steer} \in (0, \min\{\Delta^l, \underline{\sigma}\}]$ such that we have the following result.

Proposition 7 *Banning self-preferencing only has the following effects:*

	<i>M's choice of mode</i>	Π	π	CS	Δ	W
<i>if</i> $\sigma \geq \Delta^l$	<i>Seller</i>
<i>if</i> $\underline{\sigma}^{steer} < \sigma < \Delta^l$	<i>Seller</i>	↓	.	.	.	↓
<i>if</i> $\sigma \leq \underline{\sigma}^{steer}$	<i>Dual (no self-preferencing)</i>	↓	.	↑	.	↑

“.” = not changing; “↑” = increasing; “↓” = decreasing.

When $\sigma > \underline{\sigma}^{steer}$, banning self-preferencing makes the dual mode less profitable than the pure seller mode, so M switches to the seller mode, with implications that are the same as in the first part of Proposition 6.

When $\sigma \leq \underline{\sigma}^{steer}$, M continues in dual mode after the ban, and the ban restores on-platform competition and the showrooming constraint on τ . Both of these decrease the final price level and benefit consumers. Nonetheless, the innovation level remains unchanged because imitation by M implies that S still makes zero profit in equilibrium. Thus, the increase in welfare simply reflects lower prices and therefore a larger number of transactions.

□ **Ban on imitation only.** Banning imitation is equivalent to M committing not to imitate. This makes it possible for M to optimally sustain S 's innovation incentive by setting a commission that satisfies the innovation constraint $\tau \leq \bar{\tau}$ (defined in Lemma 1), as in Proposition 3. Reflecting this observation, we obtain the following post-intervention equilibrium in dual mode, which is similar to that in Proposition 3, but with a higher fee (given that self-preferencing is still allowed).

- If $b + \Delta^l \leq \bar{\tau}$ or

$$\bar{\tau}G(v + \sigma + b) \geq (b + \max\{\sigma, \Delta^l\})G(v) \quad (14)$$

holds, M sets $\tau = \bar{\tau}$ and S chooses $\Delta = \bar{\Delta}$. In this case, S sells to all consumers through the marketplace at $p_i^* = \bar{\Delta} - \sigma$.

- If $b + \Delta^l > \bar{\tau}$ and (14) does not hold, M sets $\tau = b + \Delta^l$ and S chooses $\Delta = \Delta^l$. If $\Delta^l > \sigma$, then S sells to all consumers through the marketplace at $p_i^* = \tau$. If $\Delta^l \leq \sigma$, then M sells to all consumers at $p_m^* = b + \sigma$.

Profits are $\Pi = \max\{\bar{\tau}G(v + \sigma + b), (b + \max\{\sigma, \Delta^l\})G(v)\}$ and $\pi = 0$, so that M 's profit is higher than in both pure modes. Comparing the post-intervention and the pre-intervention equilibria of the dual mode, the ban has an effect only when it induces M to change its commission to $\tau = \bar{\tau}$.

Proposition 8 *Banning imitation only has the following effects:*

	<i>M's choice of mode</i>	Π	π	CS	Δ	W
<i>if</i> $b + \Delta^l > \bar{\tau}$ and (14) does not hold	<i>Dual (no imitation)</i>
<i>if</i> $b + \Delta^l \leq \bar{\tau}$ or (14) holds	<i>Dual (no imitation)</i>	↑	.	↑	↑	↑

“.” = not changing; “↑” = increasing; “↓” = decreasing.

¹⁵There exist other equilibria involving different τ and prices, but they are all outcome-equivalent.

The proposition has interesting implications. Banning imitation is equivalent to M being able to credibly commit not to imitate, and has the effect of restoring innovation whenever M strictly benefits from the ban, i.e., when $b + \Delta^l \leq \bar{\tau}$ or (14) holds. Both of these conditions require $\bar{\tau}$ to be sufficiently large, which is equivalent to $\bar{\Delta}$ being sufficiently large relative to the innovation cost $K(\bar{\Delta})$. Thus, M benefits from a ban on imitation only when it rules out copying products with sufficiently large innovations. In practice, M could try to replicate the effect of such a policy itself by building up a reputation for only copying products based on relatively minor innovations.

Reflecting the increase in the innovation level, the ban on imitation increases both consumer surplus and welfare whenever it has an effect on M . Notice that M 's commission at $\tau = \bar{\tau}$ does not fully extract the innovation surplus given that it needs to maintain S 's innovation incentive. As a result, the higher innovation level benefits consumers, leading to larger equilibrium demand. Nonetheless, steering by M means that M 's fee is not constrained by the showrooming constraint, so the final price is still higher than in the baseline case without steering (Proposition 3).

□ **Ban on both imitation and self-preferencing.** If both imitation and self-preferencing are banned, then the dual mode equilibrium becomes equivalent to that from Proposition 3 in the baseline setup.

Proposition 9 *Banning both imitation and self-preferencing has the following effects:*

	M 's choice of mode	Π	π	CS	Δ	W
if $\sigma \geq \max\{\underline{\sigma}, \Delta^l\}$	<i>Seller</i>
if $\underline{\sigma} < \sigma < \max\{\underline{\sigma}, \Delta^l\}$	<i>Seller</i>	↓	.	.	.	↓
if $\sigma \leq \underline{\sigma}$	<i>Dual (no imitation and self-preferencing)</i>	↓	↑ if $b < \bar{\tau}$; . otherwise	↑	↑ if $b < \bar{\tau}$ or (8) holds; . otherwise	↑

“.” = not changing; “↑” = increasing; “↓” = decreasing.

We focus on the case $\sigma \leq \underline{\sigma}$, whereby M continues in dual mode after the intervention. In this case, consumer surplus and welfare increase because banning imitation and self-preferencing restores the showrooming constraint, so S 's innovation incentive is back to what it was in the baseline dual mode. Notably, banning imitation and banning self-preferencing has a synergistic effect, in the sense that the condition for Δ to increase is less restrictive in Proposition 9 than in Propositions 7 and 8. This is because the showrooming constraint on M 's fee makes it less profitable for M to set a fee that violates the innovation constraint $\tau \leq \bar{\tau}$ (which is relevant only when imitation is banned).

Finally, one could also consider the option of forcing M to share with all third-parties any proprietary data it gains from observing the sales of third-party sellers on its marketplace. This would be a way to ensure that M does not have any advantage in imitating S . However, data sharing means that now all fringe sellers can perfectly imitate S , which completely eliminates S 's innovation incentive. In our framework, it can be shown that this results in the same equilibrium outcome as the second case in Proposition 3, with M always setting $\tau = b$. Hence, the data sharing policy is equivalent to a special case of Proposition 9 where π and Δ do not increase when $\sigma \leq \underline{\sigma}$, so that the intervention is weakly worse than banning both imitation and self-preferencing.

□ **Comparing the policy approaches.** After understanding the implications of each policy option, we are now ready to compare them. We ask the following question: compared to the equilibrium induced by banning the dual mode (Proposition 5), how does using behavioural remedies instead change the

market outcome? To make the presentation concise, in the corollary below we state the changes in terms of weak inequalities.

Corollary 2 *Suppose $\sigma > 0$, so that M operates in seller mode after the dual mode is banned. Relative to banning the dual mode, behavioral remedies have the following effects:*

	<i>Remedies</i>	Π	π	CS	Δ	W
<i>if $\sigma \in (0, \underline{\sigma}^{steer}]$</i>	<i>Banning self-preferencing</i>	\uparrow	\cdot	\uparrow	\cdot	\uparrow
<i>if $\sigma > \underline{\sigma}^{steer}$</i>		\cdot	\cdot	\cdot	\cdot	\cdot
	<i>Banning imitation</i>	\uparrow	\cdot	\uparrow	\uparrow	\uparrow
<i>if $\sigma \in (0, \underline{\sigma}]$</i>	<i>Banning both</i>	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow
<i>if $\sigma > \underline{\sigma}$</i>		\cdot	\cdot	\cdot	\cdot	\cdot

“ \cdot ” = not changing; “ \uparrow ” = weakly increasing; “ \downarrow ” = weakly decreasing.

Suppose $\sigma \leq 0$, so that M operates in marketplace mode after the dual mode is banned. Relative to banning the dual mode, behavioral remedies have the following effects:

<i>Remedies</i>	Π	π	CS	Δ	W
<i>Banning self-preferencing</i>	\uparrow	\downarrow	\uparrow	\downarrow	\uparrow if (9) holds \downarrow otherwise
<i>Banning imitation</i>	\uparrow	\downarrow	\uparrow	\uparrow if $b + \Delta^l \leq \bar{\tau}$ or (14) holds \downarrow otherwise	\uparrow if Δ increases \downarrow otherwise
<i>Banning both</i>	\uparrow	\downarrow	\uparrow	\uparrow if $b \leq \bar{\tau}$ or (8) holds \downarrow otherwise	\uparrow if Δ increases or (9) holds \downarrow otherwise

“ \cdot ” = not changing; “ \uparrow ” = weakly increasing; “ \downarrow ” = weakly decreasing.

Corollary 2 shows that, in our model, targeted behavioral remedies tend to generate better outcomes than the structural intervention of banning the dual mode altogether. The only exception is the empirically less likely case in which M actually has a disadvantage in selling compared to fringe sellers (i.e. $\sigma \leq 0$). In this case, forcing M to operate in marketplace mode (by banning the dual mode) may generate higher innovation, which improves welfare if innovation is sufficiently important. As explained below Proposition 4, a necessary condition for this exception is that M is able to completely squeeze S 's margin even though M has a disadvantage in providing or selling the product in question compared to the fringe suppliers.

5 Extensions

5.1 Alternative specifications of imitation and steering

In this section, we consider the same analysis as in Section 4, but extend the model in terms of how imitation and steering work in dual mode.

□ **Constrained imitation and commitment.** In practice, technological and logistical constraints may prevent M from successfully identifying and imitating S 's innovation, so that product imitation does not always occur. In Section D of the Online Appendix we analyze the case in which after S has chosen its innovation, with probability α platform M is unable to engage in product imitation, and with the remaining probability $1 - \alpha$ the platform is able to imitate S 's product if it so chooses (as in Section 4). Meanwhile, the specification of M 's ability to steer is still the same as in Section 4.

Let $\bar{\Delta}_\alpha$ denote the solution to the first-order condition $\alpha G(v + b + \sigma) = K'(\bar{\Delta}_\alpha)$ and $\bar{\tau}_\alpha$ denote the unique solution to the indifference condition $\alpha(\bar{\Delta}_\alpha - \sigma - \bar{\tau}_\alpha)G(v + b + \sigma) - K(\bar{\Delta}_\alpha) = 0$. Thus, $\bar{\Delta}_\alpha$ and $\bar{\tau}_\alpha$ are the counterparts of $\bar{\Delta}$ (1) and $\bar{\tau}$ (7) in the baseline model. To focus on the interesting case where imitation matters in equilibrium, we assume $\bar{\Delta}_\alpha > \sigma$.

In this extension, the equilibrium of the dual mode has a similar structure as the baseline dual mode described in Proposition 3. Specifically, M sets either (i) $\tau = b + \Delta^l > \bar{\tau}_\alpha$, inducing a complete price squeeze equilibrium with S choosing the lowest innovation level Δ^l , or (ii) $\tau = \bar{\tau}_\alpha$, inducing an incomplete price squeeze equilibrium with S choosing innovation level $\bar{\Delta}_\alpha \in [\Delta^l, \bar{\Delta}]$. Whenever the equilibrium with $\tau = \bar{\tau}_\alpha$ arises, the innovation level and demand are higher than in the dual mode with unconstrained imitation in Proposition 5. Thus, the range of parameter values such that the ban on the dual mode decreases innovation, consumer surplus, and welfare becomes wider, relative to the result in Proposition 6.

Finally, an alternative interpretation of this extension is that M can imperfectly commit to not imitate (so that parameter α represents M 's commitment power). Following the logic in Proposition 8, M weakly prefers to commit to not imitate, if it can (i.e., it prefers $\alpha > 0$ over $\alpha = 0$). If α is an endogenous choice rather than being an exogenous parameter, we show that $\alpha \in (0, 1)$ is weakly optimal for M . This reflects a trade-off between sustaining S 's innovation incentive and a higher likelihood of extracting S 's innovation through product imitation.¹⁶

□ **Imperfect and value-adding imitation.** In Section E of the Online Appendix we analyze the case in which the value of M 's imitation product is $v + \Delta - \epsilon$, where $\epsilon > 0$ indicates imperfect imitation and $\epsilon < 0$ indicates value-adding imitation (which may reflect that M improves upon S 's innovation or combines its scale advantage with the innovation). We focus on the interesting case where ϵ is not too large, as otherwise M 's imitation is irrelevant. We continue to allow M to steer consumers.

As already described below Proposition 5, the main result of this extension shows that the equilibrium outcome in dual mode continues to be similar to the one in Proposition 5. If $\epsilon > 0$, imitation does not occur on the equilibrium path. Intuitively, steering already allows M to fully extract the innovation surplus through a high commission, and so imperfect imitation does not strictly increase M 's profit (but the possibility of imitation off the equilibrium path still affects S 's innovation incentive). In this case, all the existing results in Sections 4.3-4.4 continue to hold. If $\epsilon < 0$, imitation occurs on the equilibrium path. Intuitively, value-adding imitation strictly increases the innovation surplus that M can extract through a high price. In this case, the range of parameter values such that the ban on the dual mode decreases welfare becomes wider, relative to the result in Proposition 6, and the extent of this effect depends on how large the value added from imitation is.

□ **Imperfect steering.** In Section F of the Online Appendix we analyze the case in which consumers differ in the information they have regarding the offerings available on M 's marketplace. A fraction $\lambda \in [0, 1]$ of consumers are “searchers” and they are aware of S 's existence as long as S is available on M 's marketplace (as in the baseline model). The remaining fraction $1 - \lambda$ of consumers are “non-searchers” and they rely on M 's recommendation to discover products, so M can steer them (as in the model in Section 4). Specifically, after all prices are set, M chooses whether or not to show S 's product to non-searchers.

For tractability, we further assume that M has commitment power in its pricing in dual mode, i.e., sets its price before all third-party sellers (including S) do. This allows us to avoid characterizing a mixed-strategy equilibrium, which can be extremely complicated given that our model has elastic consumer

¹⁶See Madsen and Vellodi (2021) for a similar trade-off in a dynamic model with uncertain demand states.

demand G . Despite the timing difference, this setup recovers the equilibrium in Proposition 5 when $\lambda = 0$. Meanwhile, the specification of product imitation by M is still the same as in Section 4.

Compared to the dual mode with perfect steering, imperfect steering leads to a few key differences in the equilibrium characterization.

From Proposition 5, consider the equilibrium with $\tau = b + \Delta^l$ and M not imitating in equilibrium. The existence of searchers means that S can potentially induce these consumers to purchase directly. This partially restores the showrooming constraint on M 's fee so that M sets $\tau^{dual} \leq b + \Delta^l(1 - \lambda)$ in equilibrium. Reflecting this logic, τ^{dual} decreases as λ increases, which benefits S . At the same time, the lower fee is passed through as a lower final price for consumers, resulting in equilibrium demand that is strictly higher than in Proposition 5. If $\lambda \rightarrow 1$, then it is as if self-preferencing is banned, so that the description we gave above Proposition 7 applies.

Now consider the equilibrium from Proposition 5 with $\tau = b$ and M imitating and setting a high exploitative price. The existence of searchers means that S can still sell to searchers whenever M sets the exploitative price and sells to non-searchers, which benefits S . At the same time, as λ increases, this slowly restores on-platform competition between M and S , so that the final consumer price decreases.

To conclude, the existence of searchers implies a lower price in dual mode, so the range of parameter values such that the ban on the dual mode decreases consumer surplus and welfare becomes wider, relative to the setting with perfect steering (Proposition 6).

5.2 Comparison with wholesaler-retailer model

At a high level, the practice of platforms selling on their own marketplace appears similar to that of some retailers (e.g. 7-Eleven, Carrefour, Costco, Home Depot, Lidl, and Tesco) that offer their own in-house brands alongside products sourced from third-party suppliers. In this section, we discuss a key difference between these two practices and why the lessons from one do not fully translate to the other (see also Johnson, 2020, who discusses additional differences).

A fundamental difference between the two structures is that in the wholesaler-retailer structure the intermediary sets retail prices for all products regardless of whether they are in-house brands or sourced from third parties. In contrast, when platforms use the dual mode, third-party sellers on the marketplace maintain control of their prices. To analyze the implications of this difference, we first lay out a model of the wholesaler-retailer structure that is analogous to our baseline platform model.

Suppose M is a retailer, and it can source products from S and fringe suppliers. In addition, M can also source from its own in-house brand. All marginal costs are zero as in the baseline model. Then, M sells all sourced products through its own (retail) channel, competing against the direct channels of S and the fringe suppliers. Thus, M is now a multi-product firm setting prices for all its products. Each of the third-party suppliers determines the wholesale price at which they supply to M , and M decides which product(s) to source for.¹⁷ Consistent with the informational assumption in our baseline model, we assume that consumers are aware of S 's product (and its direct channel) only if M sources S 's product. All other specifications are the same as in the baseline model.

Corresponding to marketplace mode, seller mode and dual mode from our baseline platform setting, here we distinguish between the “third-party product” mode (M only sources products from third-parties), the “in-house product” mode (M only sources products from its in-house brand) and the “dual-product” mode (M sources both types of products). We formally analyze this wholesaler-retailer setting in Section G of the Online Appendix and show that M weakly prefers the dual-product mode, i.e., $\tilde{\Pi}^{dual} = \max\{\tilde{\Pi}^{3rd-party}, \tilde{\Pi}^{in-house}\}$.

¹⁷In Section G of the Online Appendix, we also consider an alternative setup in which M sets the wholesale price paid by M to third-party suppliers in case it wants to source from them. Given the wholesale price offered by M , suppliers just decide whether to supply M or not. We show that the overall insights are broadly similar in this case.

The main result of this section describes the implications of banning the dual-product mode, similar to what we did in Proposition 4.

Proposition 10 *In the wholesaler-retailer setting, banning the dual-product mode has the following effects. There exists a threshold $\tilde{\sigma} > 0$ for the additional value of M 's in-house product, such that:*

	<i>M's choice of mode</i>	Π	π	CS	Δ	W
<i>if $\sigma \geq \tilde{\sigma}^{retail}$</i>	<i>In-house</i>
<i>if $\sigma \in (0, \tilde{\sigma}^{retail})$</i>	<i>In-house</i>	.	↓	↓	↓	↓
<i>if $\sigma \leq 0$</i>	<i>Third-party</i>

“.” = not changing; “↑” = increasing; “↓” = decreasing.

From Proposition 10, banning the dual-product mode has similar implications as the baseline platform setting (Proposition 4) when $\sigma > 0$, while at the same being much less consequential when $\sigma \leq 0$. This reflects that the dual-product mode is never harmful but also less likely to be beneficial, relative to the third-party product mode (which corresponds to the marketplace mode in the baseline model).

In the wholesaler-retailer setting, when $\sigma \leq 0$, the dual-product mode never jeopardizes S 's profit and its innovation incentive. This is because S always extracts some rents through the wholesale price it sets, meaning that it has the same incentive as in the third-party product mode.¹⁸ At the same time, given that M is now a multi-product firm setting prices for all its products, the mechanism of “on-platform competition” that drives the results in the platform setting is absent. As such, the dual-product mode leads to the same retail price as the third-party product mode, which reflects the main conceptual difference between the two settings.

6 Discussion and conclusion

The practice of platforms selling products or services alongside offerings from third-party sellers is increasingly widespread. Indeed, such dual mode intermediation has clear benefits when applied across different products, including: increasing the diversity of products, allowing each product to be provided by the more efficient seller (the platform or the third-parties), saving on search costs for consumers, ensuring more stable supply, internalizing cross-product spillovers in marketing and enabling the platform to have some loss-leaders. And we have shown that there are also benefits of the dual mode even when restricting attention to the same narrow product category: combining the higher-quality third-party seller's product with the platform's more efficient channel, and exerting some competitive pressure on the third-party sellers. It is therefore not surprising that there are now companies like Mirakl, which help retailers create marketplaces for third-parties to sell alongside the products already sold by the retailers. Mirakl's customers include Best Buy, Carrefour, Darty, Kroger, Urban Outfitters, and others.

However, the use of the dual mode by platforms has also raised concerns from competition authorities regarding the possibility of distorting competition to the disadvantage of third-party sellers. While such concerns are valid, a blanket ban on the dual mode (i.e. forcing platforms to choose the same mode for all products) is likely to do more harm than good. And even when considering a ban on the dual mode within a narrow product category, our analysis suggests that such a ban often benefits third-party sellers at the expense of consumer surplus or total welfare. The main reason for this is that in dual mode, the

¹⁸If the wholesale price is set by M , then it is set at S 's marginal cost, so that S has no innovation incentive in third-party product mode or dual-product mode.

presence of the platform’s products constrains the pricing of the third-party sellers on its marketplace, which benefits consumers.

Furthermore, in the most empirically relevant case that the platform has an advantage in selling over fringe sellers, we have shown that a ban on operating in dual mode leads the platform to stop operating as a marketplace in the relevant product categories and focus on selling itself, which in turn leads to lower consumer surplus and total welfare, reflecting that it results in fewer transactions and that consumers can no longer combine the third-party seller’s superior product with the convenience benefit of trading via the platform.

One may expect that these results would be overturned once we take into account the possibility that the platform can copy the third-party seller’s innovations and steer consumers to buy from itself in dual mode. To investigate this, we have focused on a setup that maximizes the potential harm of these practices by assuming that the platform can perfectly imitate the seller’s innovation and perfectly steer consumers’ purchases. Perhaps surprisingly, even in this richer setting, a ban on the dual mode is not necessarily good for consumers or welfare, mainly because once again, in the most empirically relevant scenario mentioned above, such a ban causes the platform to switch to selling itself in order to exploit its role in facilitating product discovery. As such, banning the dual mode does not restore the third-party seller’s innovation incentive or effective price competition between products, which are the main harms caused by imitation and self-preferencing. These insights remain true in the more realistic cases with imperfect imitation and imperfect steering by the platform, although each of these two extensions move the implications of banning the dual mode closer to the baseline setup.

A key message that emerges from our paper’s results is that policy interventions that target specific behaviors by the platform are preferable to an outright ban on the dual mode. Namely, we have shown that a ban on product imitation by the platform restores sellers’ incentive to innovate, while a ban on self-preferencing restores the effective price competition between products or prevents the platform from extracting excessively high commissions from third-party sellers.

Of course, a downside of these types of behavioral policy remedies (relative to a broad stroke ban on the dual mode) is that they require continued monitoring of the platform’s conduct to be effective. For example, banning imitation would be hard to implement in practice. This is despite the fact that, as shown in our paper, the platform has an incentive to commit itself not to imitate highly innovative third-party products in order to preserve their incentives to innovate, and so would potentially benefit from an appropriately implemented ban. The difficulty comes from the fact that in practice, the platform’s own employees (working in its in-house products division) may want to opportunistically make use of data from its marketplace division. Interestingly, Amazon has an internal policy forbidding the use of non-public data about specific sellers to launch its own in-house products, and yet, as noted in Mattioli (2020), there are reports of its employees violating the policy. This suggests regulators may require the relevant platforms (e.g. Amazon) to maintain a “Chinese wall” between their respective private label and marketplace divisions, with strict penalties for violations. Similarly, to prevent self-preferencing, platforms may be required to provide public APIs that allow approved outsiders (e.g. policy makers or researchers) to audit their recommendation algorithms.

In addition to the extensions discussed in the paper to deal with imperfect imitation and imperfect steering, in the Online Appendix we also consider a version of the baseline model in which we allow for some horizontal differentiation between the products. Future work could extend this analysis to take into account the possibility that M can steer consumers. Potentially, the dual mode has additional efficiency advantages in this case because the platform can provide higher quality recommendations of which product (including its own) is best suited for the consumer. On the other hand, self-preferencing could lead to additional distortionary effects if it leads to a mismatch with the consumer’s ideal product in equilibrium. Other obvious extensions we didn’t include in the paper (but were discussed in an earlier

version) were to allow for multiple products or multiple platforms. Finally, it would be interesting to use our framework to explore the implications of platforms imposing price-parity clauses in dual mode, in which they restrict third-party sellers from selling their products at lower prices on any other channel including their own websites.

7 Appendix

7.1 Proof of Proposition 1

Following the main text, we focus on $\tau \leq b$. Suppose $\Delta \leq \bar{\Delta}$. The derivative of (3) with respect to p_i has the same sign as

$$p_i - \tau - \frac{G(v+b+\Delta-p_i)}{g(v+b+\Delta-p_i)}. \quad (7.15)$$

For all $p_i \leq \tau + \Delta \leq b + \Delta$, log-concavity of G implies that (7.15) is smaller than $\Delta - \frac{G(v+b-\tau)}{g(v+b-\tau)} \leq \Delta - \frac{G(v)}{g(v)} < 0$, where the final inequality is due to (2). Thus, S optimally sets $p_i^* = \tau + \Delta$ and $p_o^* > p_i^* - b$, earning profit $\pi = \Delta G(v+b-\tau) - K(\Delta)$. Then, the optimal innovation level Δ^* solves $G'(v+b-\tau) = K'(\Delta^*)$, where $\Delta^* = \Delta^{mkt}$ when $\tau = b$ and $\Delta^* < \bar{\Delta}$. Finally, M solves

$$\max_{\tau \leq b} \tau G(v+b-\tau),$$

where (2) implies $\tau = b$ is optimal.

7.2 Proof of Proposition 2

For all $p_m \leq b + \sigma$, the derivative of M 's profit with respect to p_m has the same sign as $\frac{G(v+b+\sigma-p_m)}{g(v+b+\sigma-p_m)} - p_m > \frac{G(v)}{g(v)} - b - \sigma > 0$ given (2).

7.3 Proof of Proposition 3 (Dual mode)

We first prove Lemma 1. For all $\tau \geq 0$, regardless of whether we have the semi-seller mode equilibrium or the price-squeeze equilibrium, S 's profit (as a function of Δ) can be summarized as

$$\pi(\Delta) = \max\{(\Delta - \sigma - \tau)G(v + \sigma + b), 0\} - K(\Delta).$$

Note that $\pi(\Delta)$ is continuous. If $\tau \leq \Delta^l - \sigma$, then the \max operator is irrelevant and $\pi(\Delta)$ is concave for all Δ , so $\bar{\Delta}$ is optimal. If $\tau \geq \bar{\Delta} - \sigma$, then $\pi(\Delta) = -K(\Delta)$ for all $\Delta \leq \bar{\Delta}$ and $\pi(\Delta)$ is decreasing for all $\Delta > \bar{\Delta}$, so Δ^l is optimal. If $\bar{\tau} \in (\Delta^l - \sigma, \bar{\Delta} - \sigma)$, then $\pi(\Delta)$ has two peak points, respectively at Δ^l and $\bar{\Delta}$. Given that $\pi(\bar{\Delta})$ is decreasing in τ while $\pi(\Delta^l) = 0$ is constant, by the intermediate value theorem, the unique cutoff $\bar{\tau}$ stated in Lemma 1 exists.

For the commission setting decision, consider $\tau \in (\max\{-\sigma, 0\}, b]$. Suppose $b \leq \bar{\tau}$. For all $\tau \leq b$, Lemma 1 implies S chooses $\bar{\Delta}$ and induces a price-squeeze equilibrium with $p_m^* = 0$ and $\Pi(\tau) = \tau G(v + \sigma + b)$, so $\tau^{dual} = b$. Suppose $b > \bar{\tau}$, then

$$\Pi(\tau) = \begin{cases} \tau G(v + \sigma + b) & \text{if } \tau \in [-\sigma, \bar{\tau}] \\ (\tau + \max\{\sigma - \Delta^l, 0\})G(v + \Delta^l + b - \tau) & \text{if } \tau \in (\bar{\tau}, b] \end{cases},$$

whereby there is a discrete change at $\tau = \bar{\tau} > \sigma - \Delta^l$ (Lemma 1). Assumption (2) implies $\Pi(\tau)$ is piece-wise increasing. The choice of optimal fee is between $\tau^{dual} = \bar{\tau}$ and $\tau^{dual} = b$, as indicated in (8). In this case, if $\tau^{dual} = \bar{\tau}$, then M induces a price-squeeze equilibrium with $p_m^* = 0$. If $\tau^{dual} = b$, then M induces either a price-squeeze equilibrium (if $\sigma < \Delta^l$) or a semi-seller mode equilibrium (if $\sigma \geq \Delta^l$), with $p_m^* = \tau^{dual} - \Delta^l + \sigma$ in both equilibria.

In sum, the equilibrium profits are

$$\Pi^{dual} = \begin{cases} bG(v + \sigma + b) & \text{if } b \leq \bar{\tau} \\ \max\{\bar{\tau}G(v + \sigma + b), (b + \max\{\sigma - \Delta^l, 0\})G(v + \Delta^l)\} & \text{if } b > \bar{\tau} \end{cases} \quad (7.16)$$

and

$$\pi^{dual} = \begin{cases} (\bar{\Delta} - \sigma - b)G(v + \sigma + b) - K(\bar{\Delta}) & \text{if } b \leq \bar{\tau} \\ 0 & \text{if } b > \bar{\tau} \end{cases}. \quad (7.17)$$

Finally, we rule out $\tau \notin (\max\{-\sigma, 0\}, b]$. For $\tau > b$, S always sells through the direct channel. By the Bertrand competition logic, it is easy to verify that equilibrium profits in the stage-3 pricing subgame are

$$\begin{aligned} \Pi &= \max\{b + \sigma - \Delta, 0\}G(v + \Delta^l) \\ \pi &= \max\{0, \Delta - \sigma - b\}G(v + b + \sigma) - K(\Delta). \end{aligned}$$

The definition of $\bar{\tau}$ in Lemma 1 implies that: (i) if $b \leq \bar{\tau}$, then S sets $\Delta = \bar{\Delta} > \sigma + b$ so that $\Pi = 0$; (ii) if $b > \bar{\tau}$, then S sets $\Delta = \Delta^l < \sigma + b$ so that $\Pi = (b + \sigma - \Delta^l)G(v + \Delta^l)$. In both cases, M 's profit is no higher than (7.16). For $\tau \leq -\sigma$ (which only arises if $\sigma \leq 0$), consumers always prefer the fringe product on the marketplace over M 's offering because $v + b - \tau \geq v + b + \sigma - p_m$. The pricing subgame unfolds as in the pure marketplace mode, with $\Pi \leq bG(v)$, which is lower than (7.16).

7.4 Proof of Corollary 1

The first two points follow from direct comparisons of profit expressions, where Π^{dual} is given in (7.16). For the last point, if $\sigma \leq 0$ then $\Pi^{dual} > \Pi^{mkt} \geq \Pi^{sell}$, so it suffices to focus on $\sigma > 0$. From (7), if $\sigma \rightarrow \bar{\Delta}$, then $\bar{\tau} \rightarrow 0$ so that

$$\Pi^{dual} - \Pi^{sell} \rightarrow (b + \sigma - \Delta^l)G(v + \Delta^l) - (b + \sigma)G(v) < 0,$$

where the last inequality is due to (2). Hence, the required cutoff $\underline{\sigma} \in (0, \bar{\Delta})$ exists. To establish the uniqueness, consider any arbitrary $\sigma > 0$ such that $\Pi^{dual} - \Pi^{sell} < 0$. Suppose $\sigma < \Delta^l$, then observe that

$$\Pi^{dual} - \Pi^{sell} = \begin{cases} bG(v + \sigma + b) - (b + \sigma)G(v) & \text{if } b \leq \bar{\tau} \\ \max\{\bar{\tau}G(v + \sigma + b), bG(v + \Delta^l)\} - (b + \sigma)G(v) & \text{if } b > \bar{\tau} \end{cases}$$

is continuous in σ , and

$$\frac{d\Pi^{dual} - \Pi^{sell}}{d\sigma} = \begin{cases} bg(v + \sigma + b) - G(v) & \text{if } b \leq \bar{\tau} \text{ holds} \\ \frac{d\bar{\tau}}{d\sigma}G(v + \sigma + b) + \bar{\tau}g(v + \sigma + b) - G(v) & \text{if } b > \bar{\tau} \text{ holds and (8) holds} \\ -G(v) & \text{if } b > \bar{\tau} \text{ holds and (8) does not hold} \end{cases}. \quad (7.18)$$

In the first case of (7.18), recall $\Pi^{dual} - \Pi^{sell} < 0$ implies $bG(v + \sigma + b) < (b + \sigma)G(v)$, so that

$$\frac{d\Pi^{dual} - \Pi^{sell}}{d\sigma} < bg(v + \sigma + b) - \frac{b}{b + \sigma}G(v + \sigma + b) < 0$$

where the last inequality is due to (2). In the second case of (7.18), applying the implicit function theorem and envelope theorem on (7), and using (2), we get

$$\begin{aligned} \frac{d\bar{\tau}}{d\sigma} &= -1 + (\bar{\Delta} - \sigma - \bar{\tau}) \frac{g(v + b + \sigma)}{G(v + b + \sigma)} \\ &\leq -1 + (\bar{\Delta} - \sigma - \bar{\tau}) \frac{g(v)}{G(v)} < 0. \end{aligned}$$

Then, recall $\Pi^{dual} - \Pi^{sell} < 0$ implies $\bar{\tau}G(v + \sigma + b) < (b + \sigma)G(v)$, so that

$$\begin{aligned} \frac{d\Pi^{dual} - \Pi^{sell}}{d\sigma} &< \bar{\tau}g(v + \sigma + b) - G(v) \\ &< \bar{\tau}g(v + \sigma + b) - \frac{\bar{\tau}}{b + \sigma}G(v + \sigma + b) < 0, \end{aligned}$$

where the last inequality is due to (2). Hence, $\Pi^{dual} - \Pi^{sell}$ always crosses zero from above and so crosses zero at most once. A similar argument applies for $\sigma \geq \Delta^l$ and is omitted here.

7.5 Proof of Proposition 4

Given equilibrium pricing, CS^{dual} is either $\int_0^\infty \max\{v + \Delta^l, v_o\}dG(v_o)$ or $\int_0^\infty \max\{v + b + \sigma, v_o\}dG(v_o)$, so higher than $CS^{mkt} = CS^{sell} = \int_0^\infty \max\{v, v_o\}dG(v_o)$. For S 's profit, the result follows immediately from comparing π^{dual} in (7.17) with $\pi^{sell} = 0$ and

$$\begin{aligned}\pi^{market} &= \max_{\Delta} \{\Delta G(v) - K(\Delta)\} \\ &> \max_{\Delta} \{(\Delta - \sigma - b)G(v + \sigma + b) - K(\Delta)\} \\ &\geq \pi^{dual},\end{aligned}$$

where the inequalities are due to the envelope theorem and (2). For innovation, if $b \leq \bar{\tau}$ or (8) holds, then $\Delta^{dual} = \bar{\Delta} > \Delta^{mkt} > \Delta^{sell} = \Delta^l$. Otherwise, $\Delta^{dual} = \Delta^{sell} = \Delta^l < \Delta^{mkt}$. Finally, if $\sigma \in (0, \underline{\sigma})$, $W^{dual} > W^{sell}$ is obvious given the results in Π , π , and CS . Suppose $\sigma \leq 0$. If $\Delta^{dual} = \bar{\Delta}$, then W decreases because $CS^{mkt} < CS^{dual}$ and

$$\begin{aligned}\Pi^{mkt} + \pi^{mkt} &= \max_{\Delta} \{(b + \Delta)G(v) - K(\Delta)\} \\ &< \max_{\Delta} \{(b + \Delta)G(v + \sigma + b) - K(\Delta)\} \\ &= \Pi^{dual} + \pi^{dual},\end{aligned}$$

by the envelope theorem. If $\Delta^{dual} = \Delta^l$, then W increases if and only if $W^{mkt} > bG(v + \Delta^l) + \int_0^\infty \max\{v + \Delta^l, v_o\}dG(v_o) = W^{dual}$, which can be simplified to (9).

7.6 Proof of Proposition 5 (Dual mode)

Using the same reasoning as in Proposition 3, we can focus on $\tau > \max\{-\sigma, 0\}$.

Stage-3 pricing subgame without imitation. Define p_i^{show} as the solution to

$$\tau G(v + \Delta + b - p_i^{show}) = \Pi_{no-imi}^{exploit}, \quad (7.19)$$

where $\Pi_{no-imi}^{exploit}$ is defined in (13). It indicates the highest inside price that S can set such that M still prefers showing S 's product instead of not showing it. The two equilibria can be formally stated as:

- *Exploitative equilibrium (without imitation).* M does not show S and M sells to all consumers, with prices $p_m^* = \min\{\tau, b\} + \sigma$, $p_i^* = \tau$, and $p_o^* \geq p_i^* - b$. Profits are

$$\Pi_{no-imi}^{exploit} = (\min\{\tau, b\} + \sigma)G(v + b - \min\{\tau, b\})$$

and $\pi = -K(\Delta)$. The equilibrium exists if and only if $p_i^{show} \leq \tau$ or $\sigma \geq \Delta$ or $\tau > b + \Delta$. Otherwise, S can deviate to $p_i = \min\{p_i^{show}, b + \Delta\} > \tau$ to be shown and it can sell to all consumers, earning a strictly positive revenue.

- *Price squeeze equilibrium (without imitation).* The construction of this equilibrium is the same as in the baseline dual mode in Section 3.3 except that we need an additional requirement $p_i^* \leq p_i^{show}$ (given (7.19)) as otherwise M can deviate by setting $p_m = \min\{\tau, b\} + \sigma$, not showing S , and earning $\Pi_{no-imi}^{exploit}$. Thus, any price profile satisfying $p_i^* = \min\{p_i^{show}, p_m^* + \Delta - \sigma\}$, $p_o^* \geq p_i^* - b$, and $p_m^* \in [\max\{\tau - \Delta + \sigma, 0\}, \min\{\tau, \tau + \sigma, b + \sigma\}]$ is an equilibrium. Our equilibrium selection rule selects the lowest p_m^* , so

$$p_i^* = \min\{p_i^{show}, \max\{\tau, \Delta - \sigma\}\}. \quad (7.20)$$

The profits are

$$\Pi_{no-imi}^{sqz} = \max\left\{\tau G(v + \sigma + b - \max\{\tau - \Delta + \sigma, 0\}), \Pi_{no-imi}^{exploit}\right\}$$

and $\pi = (p_i^* - \tau)G(v + \Delta + b - p_i^*) - K(\Delta)$. Note $\Pi_{no-imi}^{sqz} = \Pi_{no-imi}^{exploit}$ if and only if $p_i^* = p_i^{show}$. Clearly, the equilibrium exists if and only if $p_i^{show} \geq \tau$ and $\sigma < \Delta$ and $\tau \leq b + \Delta$.

Stage-3 pricing subgame with imitation. The subgame is relevant only if $\Delta > \sigma$. The two equilibria can be formally stated as:

- *Exploitative equilibrium (with imitation).* M sells to all consumers, with prices $p_m^* = \min\{\tau, b\} + \Delta$, $p_i^* = \tau$, and $p_o^* \geq p_i^* - b$. Profits are

$$\tilde{\Pi}_{imi}^{exploit} = (\min\{\tau, b\} + \Delta)G(v + b - \min\{\tau, b\})$$

and $\pi = -K(\Delta)$. This equilibrium always exists regardless of τ .

- *Price squeeze equilibrium (with imitation).* S sells to all consumers, with prices $p_i^* = \tau$, $p_o^* \geq p_i^* - b$, and $p_m^* = \tau$. Profits are

$$\Pi_{imi}^{sqz} = \tau G(v + b + \Delta - \tau)$$

and $\pi = -K(\Delta)$. This equilibrium exists if and only if $\tau = b + \Delta$. For all $\tau < b + \Delta$, M can deviate by setting the exploitative price, earning $\tilde{\Pi}_{imi}^{exploit}$. For all $\tau > b + \Delta$, S makes no sales because $p_i^* > b + \Delta$ implies that consumers buy from fringe sellers directly.

Stage-2 innovation and imitation decisions. If S sets $\Delta \leq \sigma$, then M has no incentive to imitate and we necessarily have an exploitative equilibrium without imitation. If S sets $\Delta > \sigma$, then it is straightforward to see that M strictly prefers imitating, except when $\tau = b + \Delta$ where M is exactly indifferent between imitating and not imitating. In all cases, S 's profit is always $-K(\Delta)$, so it always chooses Δ^l .

Stage-1 fee setting. Given that S always sets $\Delta^{dual} = \Delta^l$, the choice of fee does not affect Δ . If $\sigma \geq \Delta^l$, M optimally sets $\tau = b + \Delta^l$ to induce the exploitative equilibrium without imitation. If $\sigma < \Delta^l$, then M achieves the profit $\Pi^{dual} = (b + \Delta^l)G(v)$ through one of the stated strategies in the proposition.

7.7 Proof of Proposition 6

From Proposition 5, $\Pi^{dual} = (b + \max\{\sigma, \Delta^l\})G(v)$ is higher than in both pure modes, and CS is the same across all three modes (equilibrium demand is always $G(v)$).

After the ban, for $\sigma > 0$, W decreases because Π decreases (it does not change if $\sigma \geq \Delta^l$) while $\pi^{dual} = \pi^{sell} = 0$ and $CS^{dual} = CS^{sell}$. For $\sigma \leq 0$, π increases from zero to $\Delta^{mkt}G(v) - K(\Delta^{mkt}) > 0$; Δ increases from Δ^l to Δ^{mkt} ; W increases from

$$W^{dual} = (b + \Delta^l)G(v) + \int_0^\infty \max\{v, v_o\}dG(v_o)$$

to $W^{mkt} = \max_\Delta \{(b + \Delta)G(v) - K(\Delta)\} + \int_0^\infty \max\{v, v_o\}dG(v_o)$ due to the higher Δ .

7.8 Proof of Proposition 7 (Ban self-preferencing)

We first derive the overall equilibrium of the dual mode after self-preferencing is banned. Without self-preferencing, the showrooming constraint implies that M never sets $\tau > b$, so we can focus on $\tau \leq b$. Consider the stage-3 pricing subgame. If M does not imitate, the pricing subgame unfolds as in the baseline dual mode, where

$$\Pi = \begin{cases} \tau G(v + \sigma + b - \max\{\tau - \Delta + \sigma, 0\}) & \sigma < \Delta \\ (\tau + \sigma - \Delta)G(v + \Delta + b - \tau) & \sigma \geq \Delta \end{cases}$$

If M has imitated (which happens only if $\sigma < \Delta$), the standard logic of homogenous good Bertrand competition implies that the equilibrium prices are $p_i^* = p_m^* = \tau$ and $p_o^* = 0$. Consumers buy from either M or S 's product on the marketplace. The profits are $\Pi = \tau G(v + \Delta + b - \tau)$ and $\pi = -K(\Delta)$.

Comparing M 's profit with and without imitation, for each given Δ it is clear that M strictly prefers imitating if and only if $\tau - \Delta + \sigma < 0$, and does not imitate otherwise. In both cases, it is easy to check that S 's profit is always $-K(\Delta)$, so S always chooses Δ^l . In stage 1, $\tau = b$ is clearly optimal. Summarizing,

Lemma 2 (*Dual mode equilibrium with product imitation only*). M sets $\tau = b$, S participates and sets Δ^l .

- If $\sigma < \Delta^l$, then M imitates if $b - \Delta + \sigma < 0$ and does not imitate otherwise. Prices are $p_i^* = \tau$, $p_o^* = 0$, and $p_m^* = \tau$, and S sells to all consumers.
- If $\sigma \geq \Delta^l$, then M does not imitate. Prices are $p_i^* = \tau$, $p_o^* = 0$, and $p_m^* = b + \sigma - \Delta^l$, and M sells to all consumers.

Profits are $\Pi = (b + \max\{\sigma - \Delta^l, 0\})G(v + \Delta^l)$ and $\pi = 0$.

We are now ready to prove Proposition 7. Comparing $\Pi_{no-ster}^{dual} = (b + \max\{\sigma - \Delta^l, 0\})G(v + \Delta^l)$ and $\Pi^{sell} = (b + \sigma)G(v)$, notice that if $\sigma = \Delta^l$ then

$$\Pi_{no-ster}^{dual} = bG(v + \Delta^l) < (b + \Delta^l)G(v) = \Pi^{sell}$$

given (2). The existence of the unique cutoff $\underline{\sigma}^{ster} < \Delta^l$ follows from the intermediate value theorem. If $\sigma > \underline{\sigma}^{ster}$, we compare the pre-intervention dual mode with the pure seller mode, which is the same comparison as in the first part of Proposition 6. If $\sigma \leq \underline{\sigma}^{ster}$, we compare dual modes with and without self-preferencing. CS increases because equilibrium demand increases from $G(v)$ to $G(v + \Delta^l)$, which also results in higher W .

7.9 Proof of Proposition 8 (Ban imitation)

We first derive the overall equilibrium of the dual mode after imitation is banned. We have the following technical claim:

Claim 1 Consider p_i^{show} defined in (7.19), then

- $dp_i^{show}/d\Delta = 1$ so that $p_i^{show} - \Delta$ is independent of Δ ;
- if $\sigma \geq 0$, then $dp_i^{show}/d\tau > 0$; if $\sigma < 0$, then $p_i^{show} - \Delta \geq -\sigma$.
- if $\tau \geq b + \sigma$, then $p_i^{show} - \Delta > b$.

Proof. Applying the implicit function theorem on (7.19) yields $dp_i^{show}/d\Delta = 1$. If $\sigma \geq 0$, then

$$\frac{dp_i^{show}}{d\tau} = \begin{cases} \frac{G(v+\Delta+b-p_i^{show})-G(v+b-\tau)+(\tau+\sigma)g(v+b-\tau)}{\tau g(v+\Delta+b-p_i^{show})} > 0 & \text{if } \tau \leq b \\ \frac{G(v+\Delta+b-p_i^{show})}{\tau g(v+\Delta+b-p_i^{show})} > 0 & \text{if } \tau > b \end{cases},$$

where the first inequality is due to $G(v + \Delta + b - p_i^{show}) \geq G(v + b - \tau)$ as implied by the definition of p_i^{show} and $\tau \leq b \leq b + \sigma$. If $\sigma < 0$, then for all $p_i < \Delta - \sigma$, we have

$$\begin{aligned} \tau G(v + \Delta + b - p_i) &> \tau G(v + b + \sigma) \\ &> (\tau + \sigma)G(v + b) \geq \Pi_{no-imi}^{exploit}, \end{aligned}$$

implying $p_i^{show} \geq \Delta - \sigma$ by (7.19). Next, if $\tau \geq b + \Delta^l$, then for all $p_i \leq b + \Delta$, we have

$$\tau G(v + \Delta + b - p_i) > (b + \sigma)G(v) = \Pi_{no-imi}^{exploit},$$

so that $p_i^{show} > b + \Delta$ by (7.19). ■

Consider the innovation and imitation decisions in stage 2. Recall that the equilibrium in the stage-3 pricing subgame (without imitation) is described in the proof of Proposition 5, where p_i^{show} defined in (7.19) plays an

important role in determining which equilibrium prevails in the pricing subgame. To make the dependency of p_i^{show} on Δ explicit, we write $p_i^{show} = p_i^{show}(\Delta)$.

Suppose τ is such that $p_i^{show}(\Delta) - \Delta < -\sigma$ for all Δ (by Claim 1, this condition is independent of Δ). Then the constraint $p_i^* \leq p_i^{show}(\Delta)$ in (7.20) always binds in the price squeeze equilibrium whenever it arises. Taking into account both types of equilibria,

$$\pi(\Delta) = \max \left\{ p_i^{show}(\Delta) - \tau, 0 \right\} G(v + b + \Delta - p_i^{show}(\Delta)) - K(\Delta).$$

Regardless of S 's choice of Δ , M 's profit is always $\Pi = \Pi_{no-imi}^{exploit}$.

Suppose τ is such that $p_i^{show}(\Delta) - \Delta \geq -\sigma$. Then the constraint $p_i^* \leq p_i^{show}(\Delta)$ in (7.20) never binds. And

$$\pi(\Delta) = \begin{cases} -K(\Delta) & \text{if } p_i^{show}(\Delta) < \tau \\ \max \{ \Delta - \sigma - \tau, 0 \} G(v + b + \sigma) - K(\Delta) & \text{if } p_i^{show}(\Delta) \geq \tau \end{cases}. \quad (7.21)$$

Suppose, for the moment, we ignore the constraint $p_i^{show}(\Delta) \geq \tau$ in (7.21). Then, Lemma 1 implies that S optimally chooses $\bar{\Delta} > \tau + \sigma$ if $\tau \leq \bar{\tau}$ and chooses Δ^l if $\tau > \bar{\tau}$. The range of τ we are focusing on implies that $p_i^{show}(\bar{\Delta}) \geq \bar{\Delta} - \sigma$, so the upperbound on $\bar{\tau}$ in Lemma 1 implies $p_i^{show}(\bar{\Delta}) \geq \bar{\tau} \geq \tau$, satisfying the constraint in (7.21). Thus, for $\tau \leq \bar{\tau}$, S chooses $\bar{\Delta}$, resulting in $\Pi = \tau G(v + \sigma + b)$. For $\tau > \bar{\tau}$, S chooses $\Delta^l < \tau + \sigma$, so that M earns either the exploitative equilibrium profit or the price squeeze profit, depending on whether the condition for a price squeeze equilibrium (without imitation) to exist holds, i.e., $p_i^{show} \geq \tau$, $\sigma < \Delta^l$, and $\tau \leq b + \Delta^l$. In summary, M 's profit is

$$\Pi_{\tau \leq \bar{\tau}} = \tau G(v + \sigma + b) \quad (7.22)$$

$$\Pi_{\tau > \bar{\tau}} = \begin{cases} \max \left\{ \tau G(v + \Delta^l + b - \tau), \Pi_{no-imi}^{exploit} \right\} & \text{if } \tau \leq b + \Delta^l \text{ and } \sigma < \Delta^l \\ \Pi_{no-imi}^{exploit} & \text{if } \tau > b + \Delta^l \text{ or } \sigma \geq \Delta^l \end{cases}. \quad (7.23)$$

Then,

Lemma 3 (*Dual mode equilibrium with self-preferencing only*)

- If $b + \max\{\sigma, \Delta^l\} \leq \bar{\tau}$ holds or

$$\bar{\tau} G(v + \sigma + b) \geq (b + \max\{\sigma, \Delta^l\}) G(v),$$

as stated in (14), holds, then M sets $\tau^{dual} = \bar{\tau}$ and S participates and sets $\bar{\Delta}$. S sells to all consumers exclusively through the marketplace and the prices are $p_i^* = \bar{\Delta} - \sigma$, $p_o^* \geq p_i^* - b$, and $p_m^* = 0$.

- If $b + \max\{\sigma, \Delta^l\} > \bar{\tau}$ holds and (14) does not hold, then M sets $\tau^{dual} = b + \Delta^l$ and S sets Δ^l . If $\Delta^l > \sigma$, then S sells to all consumers exclusively through the marketplace and the prices are $p_i^* = \tau^{dual}$, $p_o^* \geq p_i^* - b$, and $p_m^* = b - \Delta^l + \sigma$. If $\Delta^l \leq \sigma$, then M sells to all consumers and the prices are $p_m^* = b + \sigma$, $p_i^* = \tau^{dual}$, and $p_o^* \geq p_i^* - b$.

The profits are $\Pi_{no-imi}^{dual} = \max \{ \bar{\tau} G(v + \sigma + b), (b + \max\{\sigma, \Delta^l\}) G(v) \}$ and $\pi^{dual} = 0$.

Proof. Case 1 ($b + \max\{\sigma, \Delta^l\} \leq \bar{\tau}$). We want to prove that $\tau^{dual} = \bar{\tau} \geq b + \max\{\sigma, \Delta^l\}$. By Claim 1, we know that for all $\tau \geq b + \max\{\sigma, \Delta^l\}$, we have $p_i^{show} - \Delta > b \geq -\sigma$ so that (7.22) and (7.23) apply. Hence, at $\tau = \bar{\tau}$, S chooses $\bar{\Delta}$ and M earns (7.22), that is,

$$\Pi_{\tau=\bar{\tau}} = \bar{\tau} G(v + \sigma + b). \quad (7.24)$$

- For all $\tau \geq b + \max\{\sigma, \Delta^l\}$, such that $\tau \neq \bar{\tau}$: (i) if $\tau \in [b + \max\{\sigma, \Delta^l\}, \bar{\tau})$, then $\Pi = \tau G(v + \sigma + b)$; (ii) if $\tau > \bar{\tau} \geq b + \max\{\sigma, \Delta^l\}$, then $\Pi = \Pi_{no-imi}^{exploit} = (b + \sigma) G(v)$. Neither of these profits is higher than (7.24).
- For all $\tau < b + \max\{\sigma, \Delta^l\} \leq \bar{\tau}$, either (i) $\Pi = \Pi_{no-imi}^{exploit} \leq (b + \sigma) G(v)$ (if $p_i^{show} - \Delta < -\sigma$) or (ii) $\Pi = \tau G(v + \sigma + b)$ (if $p_i^{show} - \Delta \geq -\sigma$). Both of these are lower than (7.24).

Case 2 ($b + \max\{\sigma, \Delta^l\} \geq \bar{\tau}$). We want to prove that either $\tau^{dual} = b + \max\{\sigma, \Delta^l\}$ or $\tau^{dual} = \bar{\tau}$. There are several sub-cases:

- For all $\tau \geq b + \max\{\sigma, \Delta^l\} > \bar{\tau}$, Claim 1 implies $p_i^{show} - \Delta > b \geq -\sigma$, so that the profit expression (7.23) applies. At $\tau = b + \max\{\sigma, \Delta^l\}$, simplifying (7.23) gives

$$\Pi_{\tau=b+\max\{\sigma, \Delta^l\}} = (b + \max\{\sigma, \Delta^l\})G(v), \quad (7.25)$$

while $\Pi_{\tau > b + \max\{\sigma, \Delta^l\}} = \Pi_{no-imi}^{exploit} = (b + \sigma)G(v)$, which is no higher than (7.25).

- For all $\tau \in (\bar{\tau}, b + \max\{\sigma, \Delta^l\})$, either (i) $\Pi = \Pi_{no-imi}^{exploit}$ (if $p_i^{show} - \Delta < -\sigma$); or (ii) $\Pi = \Pi_{no-imi}^{exploit} = (b + \sigma)G(v)$ (if $p_i^{show} - \Delta \geq -\sigma$, and $\tau > b + \Delta^l$ or $\sigma \geq \Delta^l$); or (iii) $\Pi = \max\left\{\tau G(v + \Delta^l + b - \tau), \Pi_{no-imi}^{exploit}\right\}$ (if $p_i^{show} - \Delta \geq -\sigma$, and $\tau \leq b + \Delta^l$ and $\sigma < \Delta^l$). None of these is higher than (7.25).
- For all $\tau \leq \bar{\tau}$, if $\sigma < 0$, then Claim 1 implies $p_i^{show} - \Delta \geq -\sigma$ for all $\tau \leq \bar{\tau}$, so that (7.22) applies and M 's profits is $\tau G(v + \sigma + b)$, which is maximized at $\tau = \bar{\tau}$, so $\Pi_{\tau=\bar{\tau}}$ is given by (7.24).
- For all $\tau \leq \bar{\tau}$, if $\sigma \geq 0$, then Claim 1 implies $p_i^{show} - \Delta$ is increasing in τ . There are two further sub-cases to consider. If $p_i^{show} - \Delta < -\sigma$ at $\tau = \bar{\tau}$, then we know $p_i^{show} - \Delta < -\sigma$ for all $\tau < \bar{\tau}$. Thus, $\Pi = \Pi_{no-imi}^{exploit}$ in this range, which is lower than (7.25). If $p_i^{show} - \Delta \geq -\sigma$ at $\tau = \bar{\tau}$, then $\Pi_{\tau=\bar{\tau}}$ is given by (7.24). For all $\tau < \bar{\tau}$, either (i) $\Pi = \Pi_{no-imi}^{exploit} \leq (b + \sigma)G(v)$ (if $p_i^{show} - \Delta < -\sigma$) or (ii) $\Pi = \tau G(v + \sigma + b)$ (if $p_i^{show} - \Delta \geq -\sigma$). Both of these are lower than (7.24).

Comparing (7.25) and (7.24), notice that (14) implies $p_i^{show} - \Delta \geq -\sigma$ (so that the profit expression (7.22) applies) at $\tau = \bar{\tau}$. This is because for all $p_i < \Delta - \sigma$,

$$\begin{aligned} \bar{\tau}G(v + \Delta + b - p_i) &> \bar{\tau}G(v + \sigma + b) \\ &\geq (b + \max\{\sigma, \Delta^l\})G(v) > \Pi_{no-imi}^{exploit}, \end{aligned}$$

where the second inequality is due to (14). Hence, if (14) holds, M optimally sets $\tau = \bar{\tau}$ in Case 2. If (14) does not hold, M optimally sets $\tau = b$ in Case 2. Combining Cases 1 and 2 yields the lemma statement. ■

We are now ready to prove Proposition 8. Clearly, Π_{no-imi}^{dual} is higher than the profit in both pure modes, so M continues in dual mode after the ban. We now compare the market outcome in this lemma with Proposition 5. After the ban:

- if $b + \max\{\sigma, \Delta^l\} > \bar{\tau}$ and (14) does not hold, then M sets $\tau^{dual} = b + \Delta^l$ and the market outcome remains the same.
- If $b + \max\{\sigma, \Delta^l\} \leq \bar{\tau}$ holds or (14) holds, then M sets $\tau = \bar{\tau}$ and S sets $\Delta = \bar{\Delta} > \Delta^l$. Π increases from $(b + \Delta^l)G(v)$ to $\bar{\tau}G(v + \sigma + b)$; π remains at zero (given the definition of $\bar{\tau}$); CS increases from $G(v)$ to $G(v + \sigma + b)$; W increases given Π and CS increase.

7.10 Proof of Proposition 9

Following Corollary 1, if $\sigma > \underline{\sigma}$, we compare the pre-intervention dual mode with the pure seller mode, which is the same comparison as in the first part of Proposition 6. If $\sigma \leq \underline{\sigma}$, we compare the pre-intervention dual mode with the baseline model dual mode. There are two possibilities:

- If $b > \bar{\tau}$ and (8) does not hold, then in the baseline dual mode M sets $\tau = b$ and S sets Δ^l , so that profits are $\Pi = bG(v + \Delta^l)$ and $\pi = 0$. Clearly, π and Δ do not change, W and CS increase because the equilibrium demand increases from $G(v)$ to $G(v + \Delta^l)$.
- If $b \leq \bar{\tau}$ or (8) holds, then in the baseline dual mode M sets $\tau = \min\{b, \bar{\tau}\}$ and S sets $\bar{\Delta}$, so that profits are $\Pi = \tau G(v + \sigma + b)$ and $\pi^{dual} = \max\{(\bar{\Delta} - \sigma - b)G(v + \sigma + b) - K(\bar{\Delta}), 0\}$. In this case, innovation increases from Δ^l to $\bar{\Delta}$. CS increases because the equilibrium demand increases from $G(v)$ to $G(v + \Delta^l)$. W increases given Δ and equilibrium demand increase.

7.11 Proof of Corollary 2

If $\sigma > 0$, we compare the seller mode with the equilibria after the behavioral remedies. Combining Proposition 6 with Propositions 7-9, a transitivity argument yields the stated results.

If $\sigma \leq 0$, we compare the marketplace mode with the equilibria after the behavioral remedies. The result from banning both imitation and self-preferencing follows immediately from Proposition 4. Consider the remaining two remedies. The decrease in Π and the increase in π and CS are straightforward. Given that $\Delta^{mkt} > \Delta^l$, behavioral remedies lead to a higher innovation than Δ^{mkt} if and only if each of the conditions for Δ to be increasing after the behavioral remedies in Propositions 7-8 hold. Finally, the welfare expressions are:

$$\begin{aligned} W^{mkt} &= (b + \Delta^{mkt})G(v) - K(\Delta^{mkt}) + \int_0^\infty \max\{v, v_o\}dG(v_o) \\ W_{no-steer}^{dual} &= (b + \Delta^l)G(v + \Delta^l) + \int_0^\infty \max\{v + \Delta^l, v_o\}dG(v_o) \\ W_{no-imi}^{dual} &= \begin{cases} (b + \Delta^l)G(v) + \int_0^\infty \max\{v, v_o\}dG(v_o) & \text{if } \Delta_{no-imi}^{dual} = \Delta^l \\ (b + \bar{\Delta})G(v + \sigma + b) - K(\bar{\Delta}) + \int_0^\infty \max\{v + \sigma + b, v_o\}dG(v_o) & \text{if } \Delta_{no-imi}^{dual} = \bar{\Delta} \end{cases} \end{aligned}$$

Comparing W^{mkt} and $W_{no-steer}^{dual}$ yields condition (9). Finally, $W_{no-imi}^{dual} < W^{mkt}$ if $\Delta_{no-imi}^{dual} = \Delta^l$, while $W_{no-imi}^{dual} > W^{mkt}$ if $\Delta_{no-imi}^{dual} = \bar{\Delta}$ (the proof of Proposition 4).

References

- Anderson, S. P. and Ö. Bedre-Defolie (2021). Hybrid platform model. *CEPR Discussion Paper No. DP16243*.
- Brito, D., P. Pereira, and J. Varela (2012). Does vertical separation necessarily reduce quality discrimination and increase welfare? *The BE Journal of Economic Analysis and Policy* 12(1).
- Calvano, E. and M. Polo (2021). Market power, competition and innovation in digital markets: A survey. *Information Economics and Policy* 54, 100853.
- De Cornière, A. and G. Taylor (2019). A model of biased intermediation. *RAND Journal of Economics* 50(4), 854–882.
- Etro, F. (2021). Product selection in online marketplaces. *Journal of Economics & Management Strategy*.
- Farrell, J. and M. L. Katz (2000). Innovation, rent extraction, and integration in systems markets. *Journal of Industrial Economics* 48(4), 413–432.
- Gilbert, R. J. (2021). Separation: A cure for abuse of platform dominance? *Information Economics and Policy* 54, 100876.
- Hagiu, A., B. Jullien, and J. Wright (2020). Creating platforms by hosting rivals. *Management Science* 66(7), 3234–3248.
- Hagiu, A. and D. Spulber (2013). First-party content and coordination in two-sided markets. *Management Science* 59(4), 933–949.
- Hagiu, A. and J. Wright (2015a). Marketplace or reseller? *Management Science* 61(1), 184–203.
- Hagiu, A. and J. Wright (2015b). Multi-sided platforms. *International Journal of Industrial Organization* 43, 162–174.
- Hagiu, A. and J. Wright (2019). Controlling vs. enabling. *Management Science* 65(2), 577–595.

- Jiang, B., K. Jerath, and K. Srinivasan (2011). Firm strategies in the mid tail of platform-based retailing. *Marketing Science* 30(5), 757–775.
- Johnson, J. (2017). The agency model and mfn clauses. *The Review of Economic Studies* 84(3), 1151–1185.
- Johnson, J. (2020). Online marketplaces and vertical integration: Prospect for harm. *Working Paper*.
- Krämer, J. and O. Zierke (2020). Paying for prominence: the effect of sponsored rankings on the incentives to invest in the quality of free content on dominant online platforms. *SSRN Working Paper 3584371*.
- Madsen, E. and N. Vellodi (2021). Insider imitation. *Working Paper*.
- Mandy, D. M. and D. E. Sappington (2007). Incentives for sabotage in vertically related industries. *Journal of regulatory economics* 31(3), 235–260.
- Mattioli, D. (2020). Amazon scooped up data from its own sellers to launch competing products. *The Wall Street Journal, April 23*.
- Rey, P. and J. Tirole (2007). A primer on foreclosure. *Handbook of industrial organization* 3, 2145–2220.
- Wang, C. and J. Wright (2020). Search platforms: Showrooming and price parity clauses. *the RAND Journal of Economics* 51(1), 32–58.
- Zenryo, Y. (2020). Platform encroachment and own-content bias. *SSRN Working Paper 3683287*.
- Zhu, F. and Q. Liu (2018). Competing with complementors: An empirical look at amazon.com. *Strategic Management Journal* 39(10), 2618–2642.