Platforms and the exploration of new products*

Andrei Hagiu† and Julian Wright‡

November 24th, 2018

Abstract

A key strategic decision for platforms is the extent to which they should facilitate the entry of untested new products and sellers alongside established products and sellers. Exploration by current buyers helps future buyers better value these products and sellers. This externality across buyers raises the possibility that there may be too little exploration from the platform's perspective. On the other hand, when sellers have market power and so price strategically, and platforms extract a share of seller revenue (i.e., charge commissions), the direction of any such bias is not obvious. We provide a theory that predicts when the platform will prefer more, less or the same level of exploration as that induced by sellers in equilibrium. Our theory can explain why platforms will sometimes have no incentive to steer buyers one way or another, or may actually want to steer buyers towards established products and sellers.

JEL classification: D23, D86, L14

Keywords: platform design, product experimentation, information

1 Introduction

A key strategic decision for platforms like Airbnb, eBay and TaskRabbit is the extent to which they should facilitate the entry of untested new (i.e., risky) products and sellers alongside established (i.e., safe) products and sellers. The entry of such products and sellers is important for the growth of platforms because they provide new choices for buyers that can sometimes turn out to be better

*We thank Luis Cabral, Doh-shin Jeon, Gunhaeng Lee, Yi Zhu, participants at the 2017 Workshop on Economics of Platforms in Barcelona, the 2017 Asia-Pacific Industrial Organisation Conference and the 2018 Online Platform Competition Conference at the University of Florida Gainesville, and seminar participants at the Collegio Carlo Alberto in Turin, Toulouse School of Economics and the University of Colorado Boulder for helpful comments. Tat-How Teh provided excellent research assistance. We are also grateful to the Collegio Carlo Alberto in Turin for funding our visit in 2018, during which we completed the research presented in this paper.

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

‡Department of Economics, National University of Singapore, Singapore 117570, E-mail: jwright@nus.edu.sg

relative to established products and sellers. However, the exploration of these risky new products and sellers creates a public good problem between today's buyers and future buyers.

By exploring (i.e., purchasing new and untested products or services), today's buyers generate valuable information about new products and sellers, which can be conveyed to future buyers through the platforms' ratings and review systems. This suggests that platforms should be concerned that there is too little exploration of new products and new sellers by buyers, especially in settings in which buyers make one-time purchases and sellers have no credible way to signal their quality to buyers.

With this positive externality in mind, one may expect platforms to seek to induce more buyer exploration. In principle, a platform has multiple ways it could steer buyers towards new products or sellers. For example, it could make new products and sellers more prominent in search or default listings, or it could provide financial incentives for buyers to try out a new product (e.g. by offering a free gift). However, there is little evidence of platforms taking such actions.

Looking across a number of prominent online platforms for products and services, including Airbnb, Amazon's marketplace, Apple's App Store, Booking.com, Coursera, eBay, edX, Etsy, Expedia, DoorDash, Google's Play Store, Grubhub-Seamless, Rakuten, Rover, TaskRabbit, Uber Eats, and Upwork, we did not find any clear-cut evidence of steering of buyers in favor of new (or riskier) products or sellers. Indeed, for most of these platforms, the default listings presented to buyers ranked established sellers first (particularly those with high reputation scores). For example, Amazon's buybox, which determines the default seller that buyers see, is only available to established highly-ranked sellers. The only exception we found was TaskRabbit, which seems to mix some new taskers between established high-reputation taskers in its default search rankings. We also did not find any evidence of a platform giving buyers a financial incentive to try a new seller or product.

In this paper, we provide a theoretical model that addresses this puzzle by showing that when sellers set prices and have market power, under various reasonable conditions, the platform has no reason to want a higher level of exploration than that induced by sellers. This possibility takes into account that a new seller or seller of a new product has an incentive to subsidize buyers initially to support exploration, provided later they can benefit through higher prices should their product or service turn out to be popular with buyers. Perhaps surprisingly, the same logic can also apply to an established seller with market power, who can compete less aggressively for some period to allow new sellers to make sales. This is because in our model, exploration also increases the profits of the established seller, who only cares about exploration when it reveals sufficiently negative information about the new sellers, in which case the established seller can raise its prices.

Because the platform is assumed to extract a share of revenues (which is what all the platforms mentioned above do), its interests are broadly aligned with the sellers' profits in determining the level of exploration. The equilibrium level of exploration is determined by the pricing of the sellers with market power, and since they are the only sellers making positive profits in our model, this is also aligned with what the platform wants. In principle, one might think that the presence of sellers' marginal costs should drive a wedge between the platform's and sellers' preferred levels of exploration. However, provided sellers have the same marginal costs and pay the same revenue shares to the platform, it turns out that in our benchmark model the level of exploration maximizing seller revenue is the same as the one maximizing seller profits, so that the platform is happy with the equilibrium level of exploration.

We analyze how this benchmark result changes when we introduce various additional factors. We find the platform prefers less exploration than that induced by sellers in equilibrium when (i) it cares about buyer surplus, and both the new and established sellers have market power, (ii) the platform can make costly investments in expanding the period-2 buyer market, and at least one of the sellers has market power, (iii) the new sellers have lower marginal costs or pay a lower revenue share to the platform relative to established sellers, and at least one of the sellers has market power, (iv) new sellers have market power and derive private benefits outside the platform from exploration on the platform, and (v) new sellers are horizontally differentiated from established sellers, and the expected value offered by new sellers is sufficiently low relative to that offered by established sellers. On the other hand, we find the platform prefers more exploration than that induced by sellers in equilibrium when (i) the platform cares about buyer surplus, and neither new nor established sellers have market power, (ii) the new sellers have higher marginal costs or pay a higher revenue share to the platform relative to established sellers, and (iii) new sellers are horizontally differentiated from established sellers, and the expected value offered by new sellers is high enough relative to that offered by established sellers. We are thus able to provide some guidance to platforms on when they should consider steering buyers towards established or new products/sellers, or when there is no reason to steer buyers at all.

The issues surrounding the entry of new products and sellers are becoming increasingly relevant. Traditionally, learning from past buyers was via word-of-mouth, which suggests the opportunity for consumers to learn from other consumers' exploration may have been slow and limited. The availability of online user reviews and the use of recommendation algorithms to match consumers' preferences via online platforms has arguably made learning from past consumer purchases much faster and more effective. Moreover, online marketplaces have become increasingly important for influencing consumers' purchase decisions.

It is well understood that the introduction of new products is a key factor in driving innovation and economic progress (Bresnahan and Gordon, 1997). There exists a burgeoning literature spanning economics, management, and marketing that studies how new products enter markets, whether there is too little or too much new product entry from a social welfare perspective, and the dynamics of pricing and learning by consumers. A key novelty of our paper relative to this literature is that we focus on a platform which extracts commissions from sellers and study whether in equilibrium there is too much or too little exploration from the platform's point of view. Thus, we are the first to look at new product exploration in the platform literature. Moreover, we not only look at the entry of new products but also that of new sellers for a given service.

Most of the existing literature on new product exploration by buyers differs from what we do in that buyers are assumed to be long lived, so they can benefit from the results of their own exploration. This means buyers are strategic and the underlying mechanisms are quite different from ours. Examples of papers in which there is a single buyer include Bergemann and Välimäki (1996) and Krähmer (2003).

Bergemann and Välimäki (1997) introduced two-sided learning in which both buyers and a seller of a new product can learn over time about buyers' tastes for the new product, as more buyers purchase it. They show that the equilibrium market shares display excessive sales of the new product relative to the social optimum in early stages and insufficient sales later on. The related analysis in Vettas (1998) exhibits similar features. Yu et al. (2016) obtain a number of interesting new results when consumers can wait until others try the new product first, which creates a free-riding problem among consumers. Bergemann and Välimäki (2006) focus on a single monopoly seller introducing a product of uncertain quality, and analyze the interplay between the optimal dynamics of monopoly pricing and learning by consumers about the new product. Ajorlou et al. (2018), Crapis et al. (2017), and Papanastasiou and Savva (2017) are recent examples of studies that allow for social learning between consumers in similar settings.

Bergemann and Välimäki (2000) contains the closest analysis to our own, in that buyers in their model have no strategic reason to explore. Most relevant is Section 2 of their paper, where they provide a two-period example with one established and one new seller, and a unit mass of buyers. In period 2, buyers either all find out the true value of the new seller's product or none do, where the probability of finding out is increasing in the number of buyers that explore the new seller's product in period 1. Given that each individual buyer has no reason to "invest" in exploration in period 1, the new seller must subsidize exploration in period 1, as is the case in our setting. They find that in equilibrium there is excessive exploration from a social welfare perspective. This parallels our result of excessive exploration in the extension of our benchmark when the platform also cares about buyer surplus in addition to extracting a share of seller revenues *and* there is one seller of each type. The underlying logic is similar: excessive exploration is due to the fact that in period 2, both the established and the new seller benefit from exploration, whereas buyer surplus is reduced by exploration.

Our paper is distinct from Bergemann and Välimäki (2000) and all other relevant papers mentioned above in several key ways. First, we focus on a platform which extracts commissions from sellers, and look at its incentive to increase or decrease the level of exploration relative to the equilibrium level induced by sellers (as opposed to comparing the equilibrium level and the socially efficient level of exploration). Second, we conduct our analysis for four different market structures, whereas previous papers typically only focus on one market structure. Third, in our model, exploration is determined along a continuum and is typically interior in equilibrium, which is more realistic than the simplifying all-or-nothing structure used in some of the above papers. This is because we allow the new product's value to be drawn from a continuum of valuations (rather than just a high or low realization) and we use a learning technology based on the truth and noise model introduced in Lewis and Sappington (1994) and further developed by Johnson and Myatt (2006).

2 Model

We consider a model with two products (or services), a platform and two periods. In each period there is a continuum (measure one) of buyers that only live for one period and wish to buy one unit of a product, which they can only do so through the platform. Buyers have an outside option with value normalized to zero. The second period is discounted by all parties by δ.

There are two types of products. The first is an established product, which we will refer to as the "safe" product. A safe product has known value, denoted u_s. The second type of product is a new untested product, which we will refer to as the "risky" product. A risky product has uncertain value to buyers, denoted u_r. One could also interpret the safe product as corresponding to an established seller of a given service (i.e., a safe seller), while the risky product corresponds to a new untested seller of that same service (i.e., a risky seller). The safe seller is known to be able to provide the service at consistent quality, but there is uncertainty regarding the risky seller's ability.

We assume u_r is drawn from the cumulative distribution G over the support $[u_L, u_H]$ and represents buyers' true value associated with the risky product. We assume that $\overline{u}_r \equiv E[u_r] < u_s < u_H$ to make the problem interesting.[1] Otherwise, in the period-1 equilibrium, buyers would either always choose the risky product (if $u_s \leq \overline{u}_r$) or they would always choose the safe product (if $u_s \geq u_H$).

Following Lewis and Sappington (1994) and Johnson and Myatt (2006), we model learning across the two periods via a "truth-or-noise" technology. If λ period-1 buyers try the risky product (i.e., they explore), then period-2 buyers observe a signal x, which with probability $F(\lambda)$ equals the true value u_r and with probability $1 - F(\lambda)$ is an i.i.d. draw from $G[u_L, u_H]$. Note if $F(\lambda) = 1$, then in period 2, the signal will correspond to the true value u_r, and so all buyers are fully informed of the risky product's quality. $F(\lambda)$ represents the accuracy of the signal obtained by period-2 buyers. We assume $F(1) \leq 1$ and F is positive, strictly increasing and concave (i.e., $F > 0$, $F' > 0$ and $F'' < 0$). This means that when more period-1 buyers explore, period-2 buyers are more likely to learn the true value u_r (through word of mouth, feedback systems, etc.); however, there are diminishing returns to exploring in terms of improving the accuracy of the signal. Diminishing returns seems natural, and is required to ensure intermediate levels of exploration will be chosen.

Upon receipt of the signal x, a buyer is unable to distinguish between truth or noise. However, Bayesian updating implies that the buyer's posterior expectation conditional on a signal x and given a level of exploration λ is

$$E_\lambda[u_r|x] \equiv F(\lambda)\, x + (1 - F(\lambda))\, \overline{u}_r. \tag{1}$$

Given $F(\lambda) > 0$, the buyer's posterior expectation is strictly increasing in x. We adopt this particular specification for its tractability. It captures in a convenient way that, when there is more exploration, buyers place more weight on the signal generated from that exploration. For example, one could interpret the signal x as capturing the average feedback score for the risky product or seller. The weight placed on this signal as representing the true quality of the risky product or seller would then depend on the number of useful reviews, which is a function of λ, the number of buyers that try this seller's product in period 1.

[1] If buyers are risk averse, this can capture that the certainty equivalent value of u_r is lower than u_s even if the expected value of u_r is higher than u_s.

We also assume that $E_0\left[u_r|u_H\right] > u_s$, or equivalently

$$u_H > \overline{u}_r + \frac{u_s - \overline{u}_r}{F\left(0\right)}. \tag{2}$$

This requires $F\left(0\right) > 0$, which was assumed above. Nothing in our analysis hinges on this assumption: it is made for expositional convenience. Specifically, it allows us to focus on interior solutions and to avoid cumbersome expressions in integration boundaries. Assumption (2) can be interpreted to mean that by period 2, even without any exploration, buyers interpret the most positive signal possible to imply that the risky product offers higher expected utility than the safe product. Thus, the signal obtained in period 2 may reveal some useful information even if there is no buyer exploration: this may represent learning via external channels (e.g. an independent expert review). Obviously, in case u_H is infinitely high (i.e., a distribution with no upper bound on the support), even a very small amount of exogenous learning between period 1 and period 2 would be sufficient for (2) to hold.

It is important to note that there is no asymmetric information in our model. In period 1, neither buyers, sellers, nor the platform know the realization of u_r. Thus, signaling of product quality by sellers is not possible in our setup. In period 2, everyone observes the same public signal x. Thus, our model implicitly assumes there is a true quality u_r that exploration can help (imperfectly) reveal. An alternative setup would be to assume that u_r is a buyer-specific match value. Then the same analysis would apply except that it is no longer reasonable to assume that the sellers can observe the signal x which is specific to each buyer, or equivalently that they can set a price for each different value of x. In the Online Appendix A, we show the results in our benchmark setting still hold in this case. Another alternative is that u_r is known to the risky seller(s), but not to the buyers, the safe seller(s) or the platform, so that there is asymmetric information regarding u_r. This means in case a risky seller has market power, it would have the potential to signal its true quality to buyers, the safe sellers and the platform in period 2 through the level of exploration it chooses in period 1. We have chosen to abstract away from such signaling mechanisms in order to focus on the public good problem raised by exploration. This is also a reasonable assumption in many of the contexts we have in mind. For instance, when first joining Airbnb, TaskRabbit or Upwork, hosts, taskers and contractors do not actually know themselves whether they will be good at providing the corresponding services. Similarly, often sellers of new products on Amazon or eBay do not know how well their products will fit buyer tastes.

Sellers compete in prices, setting prices simultaneously at the beginning of each period. We allow each seller to set different prices to different segments of identical buyers, e.g. set a lower price to a fraction λ of buyers and a higher price for all other buyers, so as to sell to the fraction of λ buyers only. In practice, online sellers can do this with discount or promotional codes that are only issued to a fraction λ of the buyers. All sellers have marginal cost c, regardless of the product (safe or risky) that they provide. We assume $u_L \geq c$, so sales of the risky product are feasible even for the lowest possible realization of u_r. We will consider four different seller market structures:

1. Identical sellers for both the safe and the risky products.

2. One seller for the risky product and identical sellers for the safe product.

3. One seller for the safe product and identical sellers for the risky product.

4. One seller for the safe product and one seller for the risky product.

When there are identical sellers for a product (safe or risky), Bertrand competition implies that product will be priced at marginal cost in both periods. Case 1 may for instance be representative of contexts in which both the safe and risky products are commoditized and can be readily supplied by many sellers. Cases 2 and 4 may be appropriate when a seller has a patent for a new product that has not yet been market tested or when there is a new seller of a given service (the uncertainty is then related to that particular seller's ability to provide the service). Case 3, which is probably the least common in platform contexts, may nonetheless be appropriate when there is an incumbent monopoly seller of an existing known product facing entry from a potentially disruptive and competitively supplied new product that has uncertain value. For instance, this could happen to a branded product on Amazon when it starts facing competition from sellers of non-branded versions of the product, or resellers of the same branded product who sometimes turn out to offer a better buyer experience than the original seller. Finally, case 4 fits best examples of platforms for services, where each seller has their own individual ability to provide the service. With a slight abuse of language, from now on we will say "safe sellers" when referring to either established sellers of a given service or sellers of the safe product, and "risky sellers" when referring to either new sellers of a given service or sellers of the risky product.

The platform enables buyers and sellers to trade. In exchange for its services, we assume that in each period the platform extracts a share α of sellers' revenue. This is equivalent to charging a percentage commission (or proportional fee) to sellers.[2] We assume $0 < \alpha < 1$ and treat α as exogenous in our analysis, capturing factors outside of our model which pin it down (e.g. bargaining power of sellers, moral hazard on the part of sellers, competition between platforms, etc.). Importantly, we assume α is the same across periods and across sellers, and that buyers are not charged any fees, assumptions that hold in most of the platform examples mentioned in the introduction. In Section 4.1 we show how our main results change if instead the platform is able to extract a share of buyer surplus, while in Section 4.3 we show how our main results change if instead the platform can extract a different share of revenue from safe sellers than from risky sellers (so α differs).

The timing of the game we consider is as follows:

Period 1a Sellers set prices for period-1 buyers.

Period 1b Period-1 buyers make their purchase decisions, thus determining λ, the number of buyers that purchase from the risky seller(s).

[2]Settings in which sellers (rather than the intermediary) set prices to consumers, and revenue is split between sellers and the intermediary based on a fixed share, are known in the literature as the "agency model". The agency model is by the far the most common arrangement for online platforms, and applies in all the examples mentioned in the introduction.

Period 2a The signal x is realized and observed by all players.

Period 2b Sellers set prices for period-2 buyers.

Period 2c Period-2 buyers make their purchase decisions.

Finally, we assume $F'(0)$ is sufficiently high, $F'(1)$ is sufficiently close to zero, and F'' is sufficiently negative so that the following function

$$\int_{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}}^{u_H} \left(E_\lambda\left[u_r|x\right] - u_s\right) dG(x) \tag{3}$$

is concave in λ, and corner solutions are ruled out. These assumptions will ensure that the respective first-order conditions considered throughout the paper uniquely characterize optimal interior solutions.

3 Benchmark results

We analyze four different cases, depending on the market structure in the seller market. In all of these scenarios, sellers have control over their prices, though they are constrained by competition.

For each case, we are interested in seeing whether the platform would prefer more or less exploration of the risky product by buyers compared to the equilibrium outcome. To determine this we start by characterizing the equilibrium level of λ (i.e., exploration) induced by sellers. We then examine whether the platform would have preferred the sellers to induce a higher or a lower level of λ. Specifically, we determine the platform's incentive to increase or decrease λ, taking as given the sellers' prices in period 1, but allowing their prices in period 2 to reflect the off-equilibrium increase or decrease in λ. This corresponds to the direction of the strategic effect of platform steering, i.e. its effect on the level of exploration holding constant any direct effect of the steering instrument itself on payoffs (including via any change in period-1 prices). We illustrate this point by analyzing a particular model of platform steering in the Online Appendix B.

3.1 Competitive sellers of both types

In this scenario, there are identical price-setting sellers competing for buyers, both for the safe product and for the risky product. Given sellers have to share α of their revenue with the platform but face a cost of c, Bertrand-style price competition implies each product is priced at $\frac{c}{1-\alpha}$ in both periods and all sellers make zero profits in both periods. As a result, sellers of the risky product cannot benefit from exploration, and therefore they have no incentive to discount the risky product below its cost in period 1 to any buyers. Thus, absent any platform intervention and given that $u_s > \overline{u}_r$, there will be no exploration in equilibrium (i.e., $\lambda^* = 0$), so all buyers only purchase the safe product in both periods.

Let us now determine the platform's preference regarding the level of λ, taking as given that sellers price at $\frac{c}{1-\alpha}$ in period 1. Since Bertrand competition ensures the sellers set the price $\frac{c}{1-\alpha}$ in period 2 irrespective of how much exploration happens in period 1, the platform's expected profit in this case is simply

$$\alpha (1+\delta) \left(\frac{c}{1-\alpha} \right),$$

which is clearly independent of λ. We can thus state our first result.

Proposition 1 *(Competing sellers of both types) With two or more identical sellers of the safe product and two or more identical sellers of the risky product, there is no exploration in equilibrium, while the platform is indifferent over the level of exploration.*

Proposition 1 implies that even though in equilibrium there is no buyer exploration of risky products, the platform has no incentive to do anything to encourage higher (i.e., positive) levels of exploration. This reflects that competition fully pins down seller revenues, so increasing the level of exploration cannot help the platform extract more revenue from sellers.

3.2 Single seller of risky product

Suppose now there is a single seller of the risky product (the risky seller) and multiple identical sellers of the safe product (the safe sellers). Given Bertrand competition, the safe sellers always price at $\frac{c}{1-\alpha}$ in both periods. Clearly, the safe sellers make zero profit and so do not care about the level of exploration. The risky seller on the other hand, has market power, and can gain from exploration.

In period 2, the risky seller can make positive sales only when it offers higher expected surplus (i.e., $E_\lambda [u_r|x] - p_r \geq u_s - \frac{c}{1-\alpha}$, where p_r is the risky seller's price in period 2). It will therefore set $p_r = E_\lambda [u_r|x] - \left(u_s - \frac{c}{1-\alpha} \right)$ and will only want to sell in period 2 if $(1-\alpha) p_r \geq c$. This requires the expected value of the risky product $E_\lambda [u_r|x]$ to be higher than the value of the safe product u_s. Recalling that $E_\lambda [u_r|x] = F(\lambda) x + (1 - F(\lambda)) \overline{u}_r$, this condition can be written equivalently as $x \geq \overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}$. Thus, prior to observing the signal x, and taking into account that its margin in period 2 is $(1-\alpha) p_r - c = (1-\alpha)(E_\lambda [u_r|x] - u_s)$, the risky seller's period-2 expected profit is

$$\pi_r (\lambda) = (1-\alpha) \int_{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}}^{u_H} (E_\lambda [u_r|x] - u_s) \, dG(x).$$

Differentiating $\pi_r(\lambda)$ with respect to λ, we obtain

$$\pi_r' (\lambda) = (1-\alpha) F'(\lambda) \int_{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}}^{u_H} (x - \overline{u}_r) \, dG(x) > 0, \tag{4}$$

where we have used that $E_\lambda [u_r|x] - u_s = 0$ when $x = \overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}$. Thus, while the risky seller wins the market more often when signal accuracy is improved (i.e., increasing the value of F through a higher λ), this has no effect on its period-2 profits because for the marginal signal realization, the risky seller makes no profit. Nevertheless, improving signal accuracy always increases the risky seller's expected profits in period 2. Given the safe product is on average better than the risky product, the risky seller can only make profit in period 2 if it is revealed to have high quality. That is, the risky seller only sells in period 2 when the signal received by buyers is above $\overline{u}_r$ by a sufficiently large amount to make up for the higher certainty value u_s of buying from the safe sellers. In these cases, putting more weight on the signal (i.e., increasing $F(\lambda)$) increases the risky seller's profits. For this reason, from the perspective of period 2, the risky seller has an incentive to set λ as high as possible.[3]

In period 1, the risky seller can charge at most $\frac{c}{1-\alpha} - (u_s - \overline{u}_r)$ to any buyer it wants to attract. Thus, if it offers this discounted price to λ buyers, and a higher price to any other buyers,[4] the risky seller's profit in period 1 is $\lambda\left((1-\alpha)\left(\frac{c}{1-\alpha} - (u_s - \overline{u}_r)\right) - c\right)$, and its total expected profit is therefore

$$-\lambda(1-\alpha)(u_s - \overline{u}_r) + \delta\pi_r(\lambda). \tag{5}$$

Due to competitive pricing by the safe sellers, the risky seller fully internalizes both the period-2 gain in total surplus and the period-1 loss in total surplus due to exploration. As noted in the model setup, we are assuming the risky seller is able to set a discounted price for some fraction of the buyers, which is consistent with sellers offering online promotional codes that are only issued to a fraction λ of the buyers.[5]

Now compare this to the platform's objective, which is to maximize its revenue from seller commissions. Taking as given the period-1 equilibrium prices charged by each type of seller, the platform's expected profit as a function of λ is

$$\begin{aligned}&\alpha\left(\lambda\left(\frac{c}{1-\alpha} - (u_s - \overline{u}_r)\right) + (1-\lambda)\frac{c}{1-\alpha}\right)\\&+\alpha\delta\left(G\left(\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}\right)\frac{c}{1-\alpha} + \int_{\overline{u}_r + \frac{u_s-\overline{u}_r}{F(\lambda)}}^{u_H}\left(E_\lambda[u_r|x] - \left(u_s - \frac{c}{1-\alpha}\right)\right)dG(x)\right).\end{aligned}$$

Collecting the common $\frac{c}{1-\alpha}$ terms and factoring by $\frac{\alpha}{1-\alpha}$ this can be rewritten as

$$\frac{\alpha}{1-\alpha}\left((1+\delta)c - \lambda(1-\alpha)(u_s - \overline{u}_r) + \delta\pi_r(\lambda)\right). \tag{6}$$

Clearly, (6) is proportional to (5) plus a constant term with respect to λ, so the platform has no incentive to increase or decrease the level of exploration from the one induced by the risky seller, i.e. which maximizes (5).

The equilibrium level of exploration induced by the risky seller is determined by the first-order

[3]This property is reminiscent of the result in Johnson and Myatt (2006) that a monopolist prefers either the distribution of buyer willingnesses-to-pay that is most dispersed or the one that is least dispersed. Here, $\lambda = 1$ can be interpreted as the case with maximum dispersion. The reason $\lambda = 0$ is never optimal in our setting is that the risky seller faces competition from the safe sellers.

[4]In order to ensure the existence of an equilibrium in the period-1 pricing game, we assume that whenever only one type of seller has market power and buyers are indifferent between buying from the two types of sellers, they all break the tie in favor of the seller with market power.

[5]If for some reason the risky seller cannot set discriminatory prices, we can obtain the same results as above if the risky seller can instead incur a cost for each buyer that it attracts to buy its product, where the cost to attract the buyer is an increasing function of the surplus shortfall of the risky seller's product. The details are provided in Online Appendix C.

condition from maximizing (5). Specifically, using (4), it solves

$$F'(\lambda)\int_{\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}}^{u_H}(x-\overline{u}_r)\,dG(x)=\frac{u_s-\overline{u}_r}{\delta}. \tag{7}$$

The left-hand side of (7), which is proportional to $\pi_r'(\lambda)$, is always positive. Furthermore, the derivative of the left-hand side of (7) in λ, which is proportional to $\pi_r''(\lambda)$, is negative given the assumption that (3) is concave in λ. Our assumptions on $F'(0)$ and $F'(1)$ then ensure there is a unique interior solution, and so a positive level of exploration.

In summary, we have shown:

Proposition 2 (Single seller of risky product) *With two or more identical sellers of the safe product and one seller of the risky product, the platform prefers no change in exploration relative to the (positive) equilibrium level.*

To understand the result in Proposition 2, note first that under profit sharing, the risky seller's and the platform's objectives in setting λ would be the same given that the safe sellers make no profit. However, revenue sharing and profit sharing are in general different because sellers incur positive marginal costs that the platform does not internalize, leading to an upward distortion in prices under revenue sharing. Given the safe sellers make no profit in each period, the total revenue generated by safe and risky sellers in each period (which is what the platform wants to maximize) is equal to the total costs incurred by sellers in each period plus the risky seller's pre-commission profit in each period. Since the marginal cost is the same regardless of the type of seller, it is irrelevant to the choice of λ. Thus, since the risky seller's pre-commission and post-commission profit are proportional (they differ by the factor $1-\alpha$), the platform prefers not to change the level of λ from the one induced by the risky seller.

Even if the safe sellers make positive profits, the above result can still hold provided these profits are fixed at some level that is independent of λ, as would be the case in the Hotelling model of product differentiation, for example. What is more critical for Proposition 2 to hold, is that the different types of sellers incur the same costs and pay the same revenue shares to the platform. As will be shown in Section 4.3, once the platform obtains a different share of revenue across different sellers, or the two types of sellers' costs are different, the equivalence between the platform's and risky seller's preferred level of λ will break down, even in this case with one risky seller and competitive safe sellers.

That the risky seller wants a positive level of exploration reflects that its expected profits are increasing in λ in period 2. Exploration is good for the risky seller in period 2 because it raises the weight buyers put on the public signal x rather than the noise $\overline{u}_r$, which only matters to the risky seller when it wins the market in period 2, and this only happens when x is high (i.e., above $\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}$).

Our results in this section would continue to hold if we allowed for uncertain entry of one or more risky sellers in period 2. In particular, denoting by β the probability that the risky seller faces competition from other, identical risky sellers in period 2, the risky seller's expected profits for period

2 become $(1-\beta)\pi_r(\lambda)$ since with probability β the firms are Betrand competitors and obtain no profit in period 2. The platform's expected profit becomes

$$\frac{\alpha}{1-\alpha}\left((1+\delta)c-\lambda(1-\alpha)(u_s-\overline{u}_r)+\delta(1-\beta)\pi_r(\lambda)\right).$$

Note in case identical risky sellers enter in period 2, the price will be set to recover costs, so is $\frac{c}{1-\alpha}$, which explains why in period 2 the platform gets $\frac{\alpha c}{1-\alpha}$ with probability β and $\frac{\alpha}{1-\alpha}(c+\pi_r(\lambda))$ with probability $1-\beta$. Thus, the rest of the analysis would remain unchanged.

3.3 Single seller of safe product

Consider now the polar opposite case, with a single seller of the safe product (the safe seller), and many identical sellers of the risky product (the risky sellers).

Given Bertrand competition, the risky sellers always price at $\frac{c}{1-\alpha}$ in both periods, and make zero profit. For this reason, it is the safe seller that has an interest in the level of λ. To calculate the safe seller's profit, note that in period 2 the safe seller can make positive sales only if $u_s-p_s\geq E_\lambda[u_r|x]-\frac{c}{1-\alpha}$, where p_s is the safe seller's price in period 2. It will therefore set $p_s=u_s-\left(E_\lambda[u_r|x]-\frac{c}{1-\alpha}\right)$ and will only want to sell if $(1-\alpha)p_s\geq c$. This requires the value of its product u_s is higher than the expected value of the risky product $E_\lambda[u_r|x]$. This condition can be written equivalently as $x\leq\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}$. Thus, paralleling the analysis in Section 3.2, the safe seller's period-2 expected profit is

$$\pi_s(\lambda)=(1-\alpha)\int_{u_L}^{\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}}(u_s-E_\lambda[u_r|x])\,dG(x).$$

Taking the derivative of $\pi_s(\lambda)$ with respect to λ and using the identity

$$\int_{u_L}^{\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}}(\overline{u}_r-x)\,dG(x)=\int_{\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}}^{u_H}(x-\overline{u}_r)\,dG(x),$$

we obtain

$$\pi_s'(\lambda)=\pi_r'(\lambda)>0. \tag{8}$$

Thus, perhaps surprisingly, the safe seller's second period profit is increasing in λ and the derivative is the same as that of the risky seller's profit when it had market power in the previous section. To understand this, recall that increasing exploration shifts the weight buyers put on noise, which has expected value $\overline{u}_r$, to the public signal x. For the case with a single risky seller, this only mattered to the risky seller when it won sales in period 2, and that only happens when the signal is high (above) $\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}$. The gain to the single risky seller from this shift was therefore $(1-\alpha)\int_{\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}}^{u_H}(x-\overline{u}_r)\,dG(x)$. Similarly, the safe seller gains from exploration to the extent that the signal is realized below $\overline{u}_r$, since the signal then shifts down the buyers' expected value of the risky seller's product and so more weight on such a signal helps the safe seller. While signal realizations above $\overline{u}_r$ have the opposite effect, the safe seller only cares about what happens when it wins sales in period 2, i.e. when the signal realization is below $\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}$. Thus, the gain to the single safe seller in period 2 from increased exploration is given by $(1-\alpha)\int_{u_L}^{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}} (\overline{u}_r - x)\, dG(x)$, which is equal to the gain from exploration to the single risky seller in the previous case, $(1-\alpha)\int_{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}}^{u_H} (x - \overline{u}_r)\, dG(x)$.

In period 1, the safe seller can charge at most $\frac{c}{1-\alpha} + u_s - \overline{u}_r$ to any buyer it wants to attract, and if it charges any more than this, buyers will prefer to purchase from the risky seller. Thus, if it offers this price only to $1-\lambda$ buyers, and sets a higher price to any other buyers, the safe seller's profit in period 1 is $(1-\lambda)\left((1-\alpha)\left(\frac{c}{1-\alpha} + u_s - \overline{u}_r\right) - c\right)$, so its total expected profit is

$$(1-\alpha)(1-\lambda)(u_s - \overline{u}_r) + \delta \pi_s(\lambda). \tag{9}$$

Due to competitive pricing by the risky sellers, the safe seller fully internalizes both the period-2 gain in total surplus and the period-1 loss in total surplus due to exploration. Each additional buyer that purchases from one of the risky sellers (i.e., explores) involves an opportunity cost to the safe seller, given it could have made a profit of $(1-\alpha)(u_s - \overline{u}_r)$ selling to each such buyer. This explains why the safe seller still internalizes the period-1 loss from exploration.

Taking as given the period-1 equilibrium prices charged by each type of seller, the platform's expected profit is

$$\alpha\left(\lambda\frac{c}{1-\alpha} + (1-\lambda)\left(\frac{c}{1-\alpha} + u_s - \overline{u}_r\right)\right)$$
$$+\alpha\delta\left(\int_{u_L}^{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}} \left(u_s - \left(E_\lambda[u_r|x] - \frac{c}{1-\alpha}\right)\right) dG(x) + \left(1 - G\left(\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}\right)\right)\frac{c}{1-\alpha}\right).$$

Collecting the common $\frac{c}{1-\alpha}$ terms and factoring by $\frac{\alpha}{1-\alpha}$ this can be rewritten as

$$\frac{\alpha}{1-\alpha}\left((1+\delta)c + (1-\alpha)(1-\lambda)(u_s - \overline{u}_r) + \delta\pi_s(\lambda)\right). \tag{10}$$

Note (10) depends on λ in exactly the same way as does (9), other than a constant proportional factor.

Taking the derivative of (9) with respect to λ, the equilibrium level of exploration induced by the safe seller is the solution to

$$\pi_s'(\lambda) = \frac{1-\alpha}{\delta}(u_s - \overline{u}_r). \tag{11}$$

And since we have shown above that $\pi_s'(\lambda) = \pi_r'(\lambda)$, the equilibrium level of exploration here is the same as that in Section 3.2, given by (7).

In summary, we have shown:

Proposition 3 (Single seller of safe product) *With two or more identical sellers of the risky product and one seller of the safe product, the platform prefers no change in exploration relative to the (positive) equilibrium level. Moreover, the equilibrium level of exploration in this case is the same as the equilibrium level when there is a single seller of the risky product.*

The logic behind the first part of the Proposition, that the platform is happy with the equilibrium level of exploration, is the same as that for Proposition 2. The logic behind the second part of the Proposition, that the equilibrium level of exploration is the same regardless of whether there is a single seller of the risky product or a single seller of the safe product follows from two observations. The first is that cost of exploration in period 1 is the same in both cases—the opportunity cost to the safe seller from exploration (losing the margin $u_s - \overline{u}_r$ on each sale) is the same as the subsidy that the risky seller much provide to make buyers willing to explore. The second is that the expected gain from exploration in period 2 is actually the same in both cases, as discussed above.

Our results in this section would continue to hold if we allowed for uncertain entry of one or more safe sellers in period 2. The logic is the same as that given in Section 3.2 in the case with uncertain entry of risky sellers.

3.4 Single seller of each type of product

Finally, suppose there is a single seller of the safe product (the safe seller) and a single seller of the risky product (the risky seller). Recall this case is also meant to encompass examples of platforms for services, where each seller has their own individual ability to provide the service. In this case, neither seller's price is tied down to cost by Bertrand competition. Instead, in this setting the sellers are asymmetric Bertrand competitors in period 2, where the asymmetry depends on the realization of the signal, the distribution of which is endogenous. For this reason, the analysis in this case does not follow in a straightforward way from the previous cases.

In period 2, the risky seller wins all buyers if and only if the signal is such that $E_\lambda[u_r|x] \geq u_s$. In this case, the safe seller sets a price equal to $\frac{c}{1-\alpha}$, while the risky seller prices at $\frac{c}{1-\alpha} + E_\lambda[u_r|x] - u_s$. If on the other hand, the signal is such that $E_\lambda[u_r|x] \leq u_s$, then the safe seller wins all buyers: the risky seller sets a price equal to $\frac{c}{1-\alpha}$, while the safe seller prices at $\frac{c}{1-\alpha} + u_s - E_\lambda[u_r|x]$. As a result, in period 2 the risky seller's expected profits are the same as in Section 3.2 (i.e., $\pi_r(\lambda)$), while the safe seller's expected profits are the same as in Section 3.3 (i.e., $\pi_s(\lambda)$). This reflects that the losing seller is willing to price down to marginal cost, which means that for the winning seller, this case is equivalent to competing against identical sellers of the other type.

We can now characterize the equilibrium level of exploration simply using the functions $\pi_r(\lambda)$ and $\pi_s(\lambda)$ and their properties. Recall from (4) and (8) that the two functions have the same derivative, which is positive and decreasing in λ so they are both increasing and concave.

Lemma 1 *The equilibrium level λ^* of exploration with a single seller of each type of product is defined by*

$$\pi_s'(\lambda^*) + \pi_r'(\lambda^*) = \frac{1-\alpha}{\delta}(u_s - \overline{u}_r). \tag{12}$$

The condition that determines the equilibrium level of exploration (12) says that the discounted joint marginal gain from exploration derived by the two sellers in period 2 must equal the period-1 cost of exploration. To understand this, note that the equilibrium level of exploration must equate the period-2 marginal gain with the period-1 marginal cost for both the risky and the safe seller. Denote by p_r^* and p_s^* the equilibrium prices of the two sellers in period 1. For the risky seller, the marginal cost of increased exploration in period 1 is the (positive) subsidy $-((1-\alpha)p_r^* - c)$, so the equilibrium level of exploration must satisfy $-((1-\alpha)p_r^* - c) = \delta\pi_r'(\lambda^*)$. For the safe seller, the marginal cost of increased exploration in period 1 is the (positive) opportunity cost of a lost sale $(1-\alpha)p_s^* - c$, so the equilibrium level of exploration must satisfy $(1-\alpha)p_s^* - c = \delta\pi_s'(\lambda^*)$. These two conditions, combined with the requirement that the difference in period-1 prices $p_s^* - p_r^*$ must equal the difference in utilities $u_s - \overline{u}_r$ to keep buyers indifferent, imply (12) in Lemma 1.

Now we wish to determine how this equilibrium level of exploration compares to what the platform would like to choose, taking as given the prices chosen by sellers in period 1. The platform's objective function can be written as

$$\alpha\left(\lambda p_r^* + (1-\lambda)p_s^* + \frac{\delta}{1-\alpha}\left(\pi_r(\lambda) + \pi_s(\lambda) + c\right)\right). \tag{13}$$

Taking into account that $p_s^* - p_r^* = u_s - \overline{u}_r$, it is clear that the derivative of (13) with respect to λ is proportional to the equilibrium condition (12). Thus, we obtain the following result.

Proposition 4 (Single seller of each type of product) *With one seller of the safe product and one seller of the risky product, the platform prefers no change in exploration relative to the (positive) equilibrium level. Moreover, the equilibrium level of exploration in this case is strictly higher than the equilibrium level when there is a single seller of one type and competing sellers of the other type.*

As explained earlier, the equilibrium level of exploration λ^* defined by (12) must equate the discounted marginal increase in the sum of period-2 seller profits to the period-1 cost of exploration $u_s - \overline{u}_r$, so that both sellers are individually happy with λ^*. Meanwhile, since the two sellers have the same costs and pay the same revenue share to the platform, the profit extracted by the platform is proportional to the sum of seller profits in each period plus a constant that does not depend on λ, which is why the platform is happy with the equilibrium level.

Furthermore, we can compare the equilibrium level of exploration in this case, determined by (12), to the one prevailing in Sections 3.2 and 3.3, which was determined by (7) and (11). The right-hand sides of the two equations are the same, but the left-hand side in (12) is twice the one in (11). This explains the second part of Proposition 4.

Similarly to the previous two cases, in which only one type of seller had market power, our results in this section would continue to hold if we allowed for uncertain entry of one or more risky sellers in period 2. Denoting by β the probability that the risky seller faces competition from other identical risky sellers in period 2, the risky seller's expected profits for period 2 become $(1-\beta)\pi_r(\lambda)$, the safe seller's expected profit for period 2 remains unchanged at $\pi_s(\lambda)$, and the platform's expected profit becomes

$$\alpha\left(\lambda p_r^* + (1-\lambda)p_s^* + \frac{\delta}{1-\alpha}\left((1-\beta)\pi_r(\lambda) + \pi_s(\lambda) + c\right)\right).$$

Thus, the rest of the analysis remains unchanged.

4 Extensions

In this section we consider five different extensions to our benchmark model. In Section 4.1 we consider how our results change when the platform also cares about buyer surplus, so it now maximizes a weighted average of seller revenue and buyer surplus. In Section 4.2 we allow the platform to invest in marketing to attract more buyers onto the platform and we show how this changes our benchmark results. In Section 4.3 we explore the possibility that the platform extracts a different share of revenue from risky sellers than from safe sellers. At the same time, since the effects are similar, we also allow for the possibility that the risky and safe sellers face different marginal costs. In Section 4.4 we consider what happens when the risky seller (with market power) obtains additional benefits outside the platform (e.g. increased sales through other channels) from increased exploration on the platform. Finally, in Section 4.5 we allow the possibility that the safe seller and the risky seller are horizontally differentiated.

4.1 Buyer surplus

In this section we explore what happens when the platform puts a weight of $\alpha_B > 0$ on buyers' surplus from each period in its objective function (in addition to extracting the share α of seller revenues). The weight α_B could reflect that the platform is able to extract some portion of the buyers' surplus through participation fees. Alternatively, it could be a shorthand for the platform's need to attract buyers in the second period if, for example, buyers have heterogeneous participation costs. In Online Appendix D we show that if the total number of buyers who participate in period 2 depends positively on their expected surplus, Proposition 5 below holds even if $\alpha_B = 0$.

The equilibrium level of exploration induced by sellers remains unchanged in each of the four market configurations given that sellers do not care about buyer surplus, but the platform's incentive to increase or decrease λ taking as given sellers' period-1 prices potentially changes due to α_B. In the case of one risky seller and competing safe sellers, or one safe seller and competing risky sellers, our results do not change. This is because in each of these cases, expected buyer surplus in each period (from the perspective of period 1, when λ is determined) is equal to the expected surplus of the product offered by competing sellers (i.e., $u_s - \frac{c}{1-\alpha}$ or $\overline{u}_r - \frac{c}{1-\alpha}$), so does not depend on the choice of λ, which implies the platform's preference over λ does not change with α_B.

Things are different in the other two market configurations. In the case with competitive sellers of both types, buyers capture the full option value created by any exploration. Sellers still have no reason to induce exploration, since they do not capture any buyer surplus (due to Bertrand competition). On the other hand, the platform's objective function taking as given sellers' period-1 prices is now

$$\alpha_B \left(\begin{array}{c} \lambda \overline{u}_r + (1-\lambda) u_s - \frac{c}{1-\alpha} \\ +\delta \left(u_s G\left(\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}\right) + \int_{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}}^{u_H} E_\lambda [u_r | x] \, dG(x) - \frac{c}{1-\alpha} \right) \end{array} \right)$$
$$+\alpha \left(\frac{c}{1-\alpha} + \delta \frac{c}{1-\alpha} \right),$$

so the platform is no longer indifferent to the level of exploration. The first-order condition in λ turns out to be identical to (7), reflecting that it is now the buyers rather than the seller with market power that obtain the period-1 loss and period-2 gain associated with exploration. Using Proposition 2, this implies that the platform prefers a positive level of λ. Thus, in this case we find there is *insufficient* exploration from the platform's perspective.

Finally, consider the case with one risky and one safe seller. Taking as given the sellers' period-1 prices, the platform's objective function (13) must now be modified by adding the new term

$$\alpha_B \left(u_s - p_s^* + \delta \left(\int_{u_L}^{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}} E_\lambda [u_r | x] \, dG(x) + \int_{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}}^{u_H} u_s dG(x) - \frac{c}{1-\alpha} \right) \right).$$

It is easily verified that the derivative of this new term in λ is equal to $-\alpha_B \delta F'(\lambda) \int_{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}}^{u_H} (x - \overline{u}_r) \, dG(x)$, which is negative. This makes sense. Buyer surplus in period 1 is constant with respect to λ because sellers' equilibrium prices make buyers indifferent between the two products. Meanwhile, buyer surplus in the second period is equal to the expected value of the surplus offered by the *losing* seller, which is decreasing in λ for the same reason that the expected profit of the *winning* seller is increasing in λ.

Consequently, due to the new term that is decreasing in λ, the platform prefers less exploration relative to the equilibrium level as determined by (12).

The following proposition summarizes the results of this section.

Proposition 5 (Platform cares about buyer surplus)
Suppose the platform places weight $\alpha_B > 0$ on buyer surplus.
(i) With competitive sellers of both types, the platform prefers more exploration relative to the equilibrium level, which is zero.
(ii) If only one type of seller has market power, the platform prefers no change in exploration relative to the (positive) equilibrium level.
(iii) With a single seller of each type, the platform prefers less exploration relative to the (positive) equilibrium level.

Whether buyer surplus is increasing or decreasing in λ (which determines the direction of the bias in Proposition 5) depends crucially on market structure. If both types of sellers are competitive, then buyers capture the upside from exploration, which is increasing in λ. In this case there is insufficient exploration once the platform takes into account buyer surplus. If on the other hand, both types of sellers have market power, then the winning seller captures the upside from exploration, and buyers are left with the surplus offered by the losing seller, which is decreasing in λ. In this case there is excessive exploration once the platform takes into account buyer surplus.

4.2 Market expansion

Suppose the market is not of a fixed size but can expand in period 2. The case in which this arises because buyers have heterogeneous participation costs, and more buyers join in period 2 when the expected surplus is higher, has already been noted in the previous section on buyer surplus and analyzed in Online Appendix D. However, this result depended on buyers being able to observe λ before they decide whether to incur their participation cost to join the platform. If period-2 buyers cannot observe the amount of exploration that has taken place on the platform in period 1 without incurring the cost of participating on the platform, then this channel via buyer surplus would be shut down. Nevertheless, market expansion still plays a role if the platform can influence the number of new buyers participating. Specifically, in this section we explore what happens when the platform can invest in order to expand the buyer market between the two periods.

Specifically, we assume there is a measure one of buyers in period 1, but in period 2 the platform can reach $1+N$ buyers at cost $C(N)$. We assume $C(0)=0$, with $C(.)$ strictly increasing and convex. This captures the realistic feature that the platform can invest in marketing to make more buyers aware of its existence.[6] We could also allow the platform to influence the number of buyers in period 1, but this would just complicate the analysis without adding any new insights. As will become clear, the interesting and novel effects are solely driven by how the number of buyers in period 2 changes in response to the choice of λ, which is what we focus on.

In the case with competing safe and competing risky sellers, nothing changes, because seller revenues and profits continue to be pinned down by competition (all prices are equal to $\frac{c}{1-\alpha}$). Thus, there is no exploration in equilibrium and the platform remains indifferent over the level of exploration.

In the other three cases, whichever sellers have market power will now take into account that their choice of λ influences the platform's choice of N, which can benefit them in period 2. Meanwhile, when determining the platform's preferences regarding λ, we do not have to consider this effect, because the platform already chooses N optimally. As a result, the platform's ability to invest in buyer market expansion for period 2 will provide a new reason for the platform to prefer less exploration than that induced in equilibrium. The logic behind the new effect that appears is the same across all these three cases, so here we only present the analysis for the case with one risky seller and competitive safe sellers. The analysis for the other two cases is provided in Online Appendix E.

We now have to redefine the risky seller's period-2 profits as $\pi_r(\lambda, N(\lambda))$, where

$$\pi_r(\lambda, N) \equiv (1-\alpha)(1+N)\int_{\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}}^{u_H} \left(E_\lambda[u_r|x]-u_s\right)dG(x)\,, \tag{14}$$

[6] Note that in our model a buyer who is aware of the platform automatically participates because there is no cost to the buyer of participating.

and $N(\lambda)$ is the period-2 level of market expansion chosen by the platform, given the level of exploration induced by the risky seller in period 1. We assume throughout that $\pi_r(\lambda, N)$ is concave in (λ, N) and that $\pi_r(\lambda, N(\lambda))$ is concave in λ.

The risky seller faces the same period-1 pricing problem as in the benchmark setting, so its total expected profit is

$$-\lambda(1-\alpha)(u_s - \overline{u}_r) + \delta\pi_r(\lambda, N(\lambda)).$$

Taking the total derivative of this expression with respect to λ, we obtain that the equilibrium level of exploration induced by the risky seller is the solution to

$$\frac{\partial \pi_r(\lambda, N(\lambda))}{\partial \lambda} + \frac{\partial \pi_r(\lambda, N(\lambda))}{\partial N} N'(\lambda) - \frac{1-\alpha}{\delta}(u_s - \overline{u}_r) = 0. \tag{15}$$

The function $N(\lambda)$ in (15) is determined by the platform's profit maximization decision regarding N given λ. Using the same steps as before, the platform's expected profit taking the sellers' period-1 prices as given can be written as

$$\frac{\alpha}{1-\alpha}\left((1+\delta(1+N))c - \lambda(1-\alpha)(u_s - \overline{u}_r) + \delta\pi_r(\lambda, N)\right) - C(N). \tag{16}$$

Given λ, the platform chooses N to maximize (16) above, implying $N(\lambda)$ is the solution to

$$\frac{\alpha}{1-\alpha}\left(\frac{\partial \pi_r(\lambda, N)}{\partial N} + c\right) = \frac{C'(N)}{\delta}. \tag{17}$$

The equilibrium level of exploration is then determined by equation (15), where $N(\lambda)$ is defined implicitly by (17). Note that the left-hand side of (17) is increasing in λ (since the cross partial $\frac{\partial^2 \pi_r(\lambda,N)}{\partial N \partial \lambda}$ is clearly positive) and the right-hand side of (17) is increasing in N, implying $N(\lambda)$ is increasing. This makes sense: if the risky seller induces more exploration, there is more upside per period-2 buyer for the risky seller, which in turn increases the platform's surplus per period-2 buyer, because it extracts a share of the seller's revenue. The platform will therefore respond by investing more in increasing the number of period-2 buyers.

The platform's objective function is (16) and its choice of N given λ is defined by (17). Using the envelope theorem, the derivative of the platform's profit (16) with respect to λ is then proportional to

$$\frac{\partial \pi_r(\lambda, N(\lambda))}{\partial \lambda} - \frac{1-\alpha}{\delta}(u_s - \overline{u}_r). \tag{18}$$

Comparing the left-hand side of (15) with (18), the difference is the term $\frac{\partial \pi_r(\lambda,N(\lambda))}{\partial N} N'(\lambda)$, which is strictly positive given $\frac{\partial \pi_r(\lambda,N(\lambda))}{\partial N} > 0$ and $N'(\lambda) > 0$. This implies the platform prefers less exploration than the equilibrium level. A similar analysis and conclusion apply to the case with one safe seller and competitive risky sellers, as well as the case with one safe and one risky seller. Relegating the proof for these cases to Online Appendix E, the following proposition summarizes the main results for this section.

Proposition 6 (Market expansion) *Suppose the platform can make a costly investment to increase the number of period-2 buyers after observing the level of exploration determined in period 1.*
(i) With competitive sellers of both types, there is no exploration in equilibrium and the platform remains indifferent over the level of exploration.
(ii) If at least one type of seller has market power, the platform prefers less exploration relative to the equilibrium level.

The reason why market expansion leads to excessive exploration from the platform's perspective when at least one type of seller has market power is that whichever sellers have market power take into account that inducing more buyers to explore raises the platform's investment in market expansion. This increases those sellers' period-2 profits, but is not a relevant consideration for the platform since it already determines the level of market expansion optimally.

4.3 Different costs and revenue shares

We now allow the platform to get different revenue shares from the two types of sellers: α_r from risky sellers and α_s from safe sellers.[7] We think of this possibility as arising for exogenous reasons, e.g. historical commitments in the case of Rover, as explained in the previous footnote. The case $\alpha_r < \alpha_s$ may reflect that risky sellers have more bargaining power vis-a-vis the platform than safe sellers, for example because safe sellers are more competitive (our case with one risky seller and competitive safe sellers), and vice-versa if the risky sellers are more competitive. The case $\alpha_r < \alpha_s$ may also reflect that the platform must offer better terms (at least temporarily) in order to attract new sellers to try out the platform. We also allow the marginal costs of the two types of sellers to be different: c_r for risky sellers and c_s for safe sellers. The case $c_r > c_s$ may reflect that safe sellers have an advantage due to scale or experience. The case $c_r < c_s$ may reflect that risky sellers use a new, more efficient technology, or that the new product is cheaper to produce.[8]

With competing sellers of both types, risky sellers price at $\frac{c_r}{1-\alpha_r}$ in both periods, whereas safe sellers price at $\frac{c_s}{1-\alpha_s}$ in both periods. Throughout this section we assume that

$$u_s - \frac{c_s}{1-\alpha_s} > \overline{u}_r - \frac{c_r}{1-\alpha_r}. \tag{19}$$

This ensures that when both types of sellers price to recover costs in period 1, the safe sellers continue to win the market. Given (19), the equilibrium level of exploration remains $\lambda^* = 0$ in this case.

[7]Among the examples of platforms discussed in the introduction, two had some difference in fees across seller types: (i) eBay offers a 10% discount on its commission for top rated sellers that meet certain conditions, including that they offer money-back guarantees to buyers; (ii) Rover charges a 15% commission (instead of its standard 20%) to sitters/walkers that were approved before March 1, 2016.

[8]Corresponding to the assumption on (3) in the benchmark setting, we assume (20) below is concave in λ.

Meanwhile, taking the sellers' period-1 prices as given, the platform's expected profit is

$$\begin{aligned}&\lambda\frac{\alpha_r c_r}{1-\alpha_r}+(1-\lambda)\frac{\alpha_s c_s}{1-\alpha_s}\\&+\delta\left(\frac{\alpha_s c_s}{1-\alpha_s}G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}\right)+\frac{\alpha_r c_r}{1-\alpha_r}\left(1-G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}\right)\right)\right),\end{aligned}$$

which, under assumption (19), is easily seen to be increasing in λ if $\frac{\alpha_r c_r}{1-\alpha_r}>\frac{\alpha_s c_s}{1-\alpha_s}$ and decreasing in λ if $\frac{\alpha_r c_r}{1-\alpha_r}<\frac{\alpha_s c_s}{1-\alpha_s}$. Thus, the platform prefers as little exploration as possible if $\frac{\alpha_r c_r}{1-\alpha_r}<\frac{\alpha_s c_s}{1-\alpha_s}$ and as much exploration as possible if $\frac{\alpha_r c_r}{1-\alpha_r}>\frac{\alpha_s c_s}{1-\alpha_s}$. Comparing with the equilibrium outcome, there is insufficient exploration from the platform's perspective when c_r is large enough relative to c_s or α_r is large enough relative to α_s. Otherwise, the platform is happy with the equilibrium level of exploration.

We next consider the case with one risky seller and competitive safe sellers. It turns out that the condition determining whether there is excessive or insufficient exploration and the underlying logic for this case are the same as for the remaining two cases (one safe seller and competitive risky sellers, and a single seller of each type). For this reason, we only present the analysis for the case with one risky seller and competitive safe sellers here; the analysis for the other two cases is provided in Online Appendix F.

When there is one risky seller and competing safe sellers, the latter price at $\frac{c_s}{1-\alpha_s}$ in both periods. For the risky seller to make any sales in period 1, the maximum price it can charge is

$$p_r=\overline{u}_r-u_s+\frac{c_s}{1-\alpha_s},$$

which is lower than the risky seller's effective cost $\frac{c}{1-\alpha_r}$ under assumption (19). In period 2, the risky seller's profit is

$$\pi_r(\lambda)\equiv(1-\alpha_r)\int_{\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}}^{u_H}\left(E_\lambda[u_r|x]-u_s+\frac{c_s}{1-\alpha_s}-\frac{c_r}{1-\alpha_r}\right)dG(x). \tag{20}$$

Thus, the risky seller's total expected profit as a function of the level of exploration chosen is

$$\Pi_r(\lambda)\equiv-\lambda(1-\alpha_r)\left(u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}\right)+\delta\pi_r(\lambda),$$

so that the level of exploration chosen by the risky seller is determined by the first-order condition $\Pi_r'(\lambda)=0$.

Meanwhile, taking sellers' period-1 prices as given, the platform's expected profit is

$$\begin{aligned}&\alpha_r\lambda\left(\overline{u}_r-u_s+\frac{c_s}{1-\alpha_s}\right)+\alpha_s(1-\lambda)\frac{c_s}{1-\alpha_s}\\&+\delta\left(\begin{array}{c}\alpha_r\int_{\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}}^{u_H}\left(E_\lambda[u_r|x]-u_s+\frac{c_s}{1-\alpha_s}\right)dG(x)\\+\alpha_s G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}\right)\frac{c_s}{1-\alpha_s}\end{array}\right),\end{aligned}$$

which can be re-written as

$$\frac{\alpha_r}{1-\alpha_r}\Pi_r(\lambda)+\left(\frac{\alpha_r c_r}{1-\alpha_r}-\frac{\alpha_s c_s}{1-\alpha_s}\right)\left(\lambda+\delta\left(1-G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}\right)\right)\right)$$
$$+\frac{\alpha_s c_s}{1-\alpha_s}+\delta\frac{\alpha_s c_s}{1-\alpha_s}.$$

Under assumption (19), the function $\left(\frac{\alpha_r c_r}{1-\alpha_r}-\frac{\alpha_s c_s}{1-\alpha_s}\right)\left(\lambda+\delta\left(1-G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}\right)\right)\right)$ is increasing (decreasing) in λ if $\frac{\alpha_r c_r}{1-\alpha_r}>\frac{\alpha_s c_s}{1-\alpha_s}$ (if $\frac{\alpha_r c_r}{1-\alpha_r}<\frac{\alpha_s c_s}{1-\alpha_s}$), so the platform prefers more (less) exploration relative to the level that maximizes the risky seller's total expected profit $\Pi_r(\lambda)$.

To understand the condition $\frac{\alpha_r c_r}{1-\alpha_r}>\frac{\alpha_s c_s}{1-\alpha_s}$, note that for each seller, the revenue extracted by the platform can be written as the sum of a term proportional to the seller's profit and a term proportional to the seller's effective cost (i.e., cost divided by the revenue share kept by the seller). Thus, in this case, the platform extracts

$$\frac{\alpha_r}{1-\alpha_r}\Pi_r(\lambda)+\frac{\alpha_r}{1-\alpha_r}c_r\left(\lambda+\delta\left(1-G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}\right)\right)\right)$$

from the risky seller and

$$\frac{\alpha_s}{1-\alpha_s}\times 0+\frac{\alpha_s}{1-\alpha_s}c_s\left(1-\lambda+\delta G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+\frac{c_r}{1-\alpha_r}-\frac{c_s}{1-\alpha_s}}{F(\lambda)}\right)\right)$$

from safe sellers.

These expressions make it clear that shifting one period-1 buyer from the safe sellers to the risky seller also results in a positive measure of period-2 buyers shifting from the safe sellers to the risky seller. Indeed, more exploration increases the chance that the expected value of the risky seller's product in period 2 exceeds the value of the safe sellers' product. Thus, shifting one buyer from safe sellers to the risky seller in period 1 creates an additional net benefit for the platform (relative to $\Pi_r'(\lambda)$) proportional to the difference between the two sellers' effective costs adjusted by the corresponding revenue shares, i.e. $\frac{\alpha_r c_r}{1-\alpha_r}-\frac{\alpha_s c_s}{1-\alpha_s}$.

The same logic applies to the other two market configurations, analyzed in Online Appendix F. The following proposition summarizes the results of this section.

Proposition 7 (Different costs or revenue shares) *Suppose the marginal cost and revenue share extracted by the platform are (c_s, α_s) for the safe sellers and (c_r, α_r) for the risky sellers, and assumption (19) holds.*
(i) With competitive sellers of both types, the platform prefers more exploration relative to the equilibrium level if $\frac{\alpha_r c_r}{1-\alpha_r}>\frac{\alpha_s c_s}{1-\alpha_s}$ and no change in the level of exploration otherwise.
(ii) If at least one type of seller has market power, the platform prefers less (more) exploration relative to the equilibrium level if $\frac{\alpha_r c_r}{1-\alpha_r}<\frac{\alpha_s c_s}{1-\alpha_s}$ (if $\frac{\alpha_r c_r}{1-\alpha_r}>\frac{\alpha_s c_s}{1-\alpha_s}$).

In the baseline model, the effective marginal costs weighted by the corresponding revenue shares were identical for both types of sellers, so the platform's incentives with respect to λ were perfectly aligned with seller profits whenever at least one type of seller had market power. Here, the difference between $\frac{\alpha_r c_r}{1-\alpha_r}$ and $\frac{\alpha_s c_s}{1-\alpha_s}$ can lead to insufficient or excessive exploration, depending on which is higher. This result underscores the difference in incentives between the platform and the sellers due to the fact that the platform extracts a share of seller revenues rather than a share of seller profits.

4.4 Private benefits from exploration

We now consider the realistic possibility that risky sellers may derive private benefits outside the platform from exploration on the platform. For instance, many sellers on eBay also sell through other channels including their own websites or other online marketplaces, and increased sales on eBay can lead to increased sales through these alternative channels.

To capture this in a simple way, we assume these private benefits only occur when the risky sellers have market power on the platform, which in our framework only arises when there is a single risky seller (i.e., we implicitly assume the market structure on the platform also applies outside the platform). In that case, the risky seller's private benefits (e.g. sales through alternative channels) are denoted $\widehat{\pi}_r(\lambda)$. We assume $\widehat{\pi}_r(\lambda)$ is increasing in λ for the same reason that $\pi_r(\lambda)$ turned out to be increasing in λ in Section 3.2: putting more weight on the signal instead of noise only has an effect on the risky seller's profit when the signal realization regarding the value of the risky seller's product is sufficiently high.

Since there are no private benefits when there are competitive risky sellers, we focus on the cases with one risky seller. Consider first the case with competitive safe sellers. The platform's revenue is unchanged since it does not capture any share of the risky seller's private benefits. Meanwhile the risky seller's total expected profit is now augmented by the private benefit $\widehat{\pi}_r(\lambda)$. Since previously the platform preferred no change in exploration relative to the (positive) equilibrium level induced by the risky seller, the private benefit now leads the risky seller to induce excessive exploration from the platform's perspective. It is straightforward to see that the same logic applies to the case with one seller of each type. In summary:

Proposition 8 (Private benefits) *If there is a single seller of the risky product who derives private benefits that are increasing in the level of exploration, the platform prefers less exploration relative to the equilibrium level.*

4.5 Horizontal differentiation between sellers

We modify our benchmark model by introducing horizontal differentiation between safe and risky sellers, in addition to the vertical quality difference. This is a realistic feature in many contexts: new hosts on Airbnb may be located in geographic areas that are not well covered by established hosts; similarly, new restaurants on Grubhub-Seamless may specalize in a particular cuisine not offered by existing restaurants; etc.

Specifically, suppose buyers are located uniformly over the Hotelling $[0,1]$ segment in product space, with one safe seller located at 0 and one risky seller located at 1. Buyers face transport cost of ty to travel a distance y in product space to a seller. The model is otherwise the same as in our benchmark case with one risky and one safe seller. Indeed, just as in our benchmark model, we find more exploration increases the expected second-period profits of both the safe and the risky seller.

Denote by λ^* the equilibrium level of exploration. Relegating the proof to Online Appendix G, we obtain the following result.

Proposition 9 *(Horizontal differentiation) Suppose there is a risky seller and a safe seller that are horizontally differentiated. Suppose also that the market is always covered and the transportation cost t is sufficiently high. Then the equilibrium level of exploration is uniquely determined by*

$$\lambda^* = \frac{1}{2} - \frac{u_s - \overline{u}_r}{6t} + \delta \frac{F(\lambda^*) F'(\lambda^*)}{27t^2} Var(u_r).$$

The platform prefers less (respectively, more) exploration relative to the equilibrium level λ^* if and only if $\lambda^* < 1/2$ (respectively, if and only if $\lambda^* > 1/2$).

This result says that, when sellers induce less than half of buyers to explore in equilibrium, the platform will want even less exploration. Conversely, when sellers induce more than half of buyers to explore in equilibrium, then the platform will want even more exploration. Thus, relative to our benchmark model, horizontal differentiation opens up the possibility that the equilibrium level of exploration can be either insufficient or excessive from the platform's perspective (recall that in the benchmark model, the platform did not want any change to the equilibrium level of exploration). Indeed, introducing horizontal differentiation makes exploration more desirable for the platform, but it also makes it easier for the risky seller to make sales despite its inferior quality. Thus, depending on whether the risky seller is able to take advantage of horizontal differentiation more or less than the value of exploration to the platform, the equilibrium level of exploration can turn out to be excessive or insufficient. This also explains why the effect of the degree of horizontal differentiation t on the equilibrium level of exploration λ^* is ambiguous. If $u_s - \overline{u}_r$ is large enough relative to $Var(u_r)$, then more differentiation leads to more exploration. However, if $u_s - \overline{u}_r$ is small relative to $Var(u_r)$, then more differentiation leads to less exploration.

5 Managerial implications and conclusion

Our paper fills a gap in the literature which to date has not considered the role of platforms or other types of intermediaries in getting buyers to try new or risky products (or services). The main question we have studied is whether platforms want to encourage more or less exploration relative to the level induced by their sellers.

The most basic insight emerging from our analysis is that insufficient exploration is less of a problem than common intuition might suggest. As we have shown, when platforms only extract a fixed share of revenues from all of their sellers, there is broad alignment between the platform's interests and the sellers' in determining the level of exploration for new products or sellers. In fact, we have identified several important factors that tend to give rise to excessive exploration from the platform's perspective:

- The platform cares about buyer surplus, and both new and established sellers have market power. In this case, when exploration reveals positive information about the new seller/product, the surplus from exploration accrues to the new seller, whereas when exploration reveals negative information, the surplus from exploration accrues to the established seller, meaning buyers are always worse off as a result of exploration. Taking this into account leads the platform to want less exploration than that induced by sellers.

- The platform can make costly investments to increase the number of participating buyers in response to the level of exploration determined by sellers, and at least one of the sellers has market power. Because exploration increases the revenue per buyer for both the platform and the sellers with market power, both the platform and the sellers prefer greater buyer market expansion when there is more exploration. The difference is that sellers do not internalize the cost of attracting the new buyers, so they induce excessive exploration.

- The new sellers have lower marginal costs or pay a lower revenue share to the platform relative to established sellers, and at least one of the sellers has market power. Relative to the sellers' incentives, the platform wishes to shift demand towards the type of seller from which it can extract a higher share of revenues or which has a higher cost (because the platform cares about seller revenue, which is equal to seller profit plus cost).

- New sellers have market power and derive private benefits outside the platform from exploration on the platform. Since the platform does not extract any share of sellers' private benefits and these private benefits are increasing in the amount of exploration, sellers induce excessive exploration from the platform's perspective.

- New sellers are horizontally differentiated from established sellers (in addition to being vertically differentiated) and the expected value offered by new sellers is sufficiently low relative to that offered by established sellers. In this case, new sellers are able to take advantage of horizontal differentiation to achieve a level of sales (i.e. exploration) that is excessive from the platform's perspective.

When one or more of these factors are present, our analysis suggests that platforms should consider steering buyers towards established products or sellers. Such steering can take the form of making established products or sellers more prominent in listings, or providing financial incentives for buyers to stick with established products/sellers. We explored one such mechanism in the Online Appendix B. Extending our framework to incorporate other specific steering mechanisms is an important avenue for future research on this topic.

Nevertheless, we also found three factors that can make insufficient exploration a relevant concern:

- The platform cares about buyer surplus and neither new nor established sellers have market power.
- The new sellers have higher marginal costs or pay a higher revenue share to the platform relative to established sellers.
- New sellers are also horizontally differentiated from established sellers and the expected value offered by new sellers is high enough relative to that offered by established sellers.

There are of course other considerations not captured in our model that could influence whether platforms need to steer buyers towards new or established sellers/products. If sellers have capacity constraints (this is especially relevant for service platforms like TaskRabbit or Thumbtack), then platforms may want to induce more exploration if doing so encourages the entry of new sellers that can help supply the market in periods of high demand. Similarly, if increasing the number of different products that sellers compete to sell on the platform increases the number of buyers participating on the platform through cross-side network effects, then platforms may want to induce more exploration of new products so as to exploit these network effects.

In this paper we focused on platforms that allow sellers to set prices to buyers directly. It would also be interesting to study exploration incentives in the context of other types of intermediaries such as a retailer that sets prices itself. Given that upstream manufacturers set wholesale prices, exploration of risky products/sellers creates a new challenge for channel coordination.

Finally, there is also the possibility that an incumbent platform may induce more or less exploration as a barrier to entry to limit competition from rival platforms. This works if the signal generated by exploration is platform specific, which is something the incumbent can potentially control. In this context, it would be interesting to study whether the incumbent platform wants to steer buyers towards established sellers or towards new sellers in order to deter entry.

6 References

Ajorlou, Amir, Ali Jadbabaie and Ali Kakhbod (2018) "Dynamic Pricing in Social Networks: The Word-of-Mouth Effect," *Management Science*, 64(2), 971-979.

Bergemann, Dirk and Juuso Välimäki (1996) "Learning and strategic pricing," *Econometrica*, 64(5), 1125-1149.

Bergemann, Dirk and Juuso Välimäki (1997) "Market Diffusion with Two-Sided Learning," *RAND Journal of Economics*, 28(4), 773-795.

Bergemann, Dirk and Juuso Välimäki (2000) "Experimentation in Markets," *Review of Economic Studies*, 67(2), 213-234.

Bergemann, Dirk and Juuso Välimäki (2006) "Dynamic Pricing of New Experience Goods," *Journal of Political Economy*, 114(4), 713-743.

Bresnahan, Timothy F. and Robert J. Gordon (1997) *The Economics of New Goods*, University of Chicago Press.

Crapis, Davide, Bar Ifrach, Costis Maglaras and Marco Scarsini (2017) "Monopoly Pricing in the Presence of Social Learning," *Management Science*, 63(11), 3586-3608.

Johnson, Justin P. and David P. Myatt (2006) "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96(3), 756-784.

Krähmer, Daniel (2003) "Entry and experimentation in oligopolistic markets for experience goods," *International Journal of Industrial Organization*, 21(8), 1201-1213.

Lewis, Tracy R. and David E. M. Sappington (1994) "Supplying Information to Facilitate Price Discrimination," *International Economic Review*, 35(2), 309–27.

Papanastasiou, Yiangos and Nicos Savva (2017) "Dynamic Pricing in the Presence of Social Learning and Strategic Consumers," *Management Science*, 63(4), 919-939.

Vettas, Nikolaos (1998) "Demand and Supply in New Markets: Diffusion with Bilateral Learning," *RAND Journal of Economics*, 29(1), 215-233.

Yu, Man, Laurens Debo, Roman Kapuscinski (2016) "Strategic Waiting for Consumer-Generated Quality Information: Dynamic Pricing of New Experience Goods," *Management Science*, 62(2), 410-435.

7 Appendix

7.1 Proof of Lemma 1

Suppose that in equilibrium, λ period-1 buyers purchase from the risky seller and the other $1-\lambda$ period-1 buyers purchase from the safe seller. The sellers' period-1 prices must satisfy

$$p_r - p_s = \overline{u}_r - u_s, \tag{21}$$

since otherwise the λ buyers who buy from the risky seller would rather buy from the safe seller, or the $1-\lambda$ buyers who buy from the safe seller would rather buy from the risky seller.[9]

While period-1 buyers are indifferent between the risky and the safe product at the equilibrium prices, the risky and the safe sellers are not indifferent over the resulting λ because they take into account their period-2 profits. Specifically, the risky seller's total discounted profits are

$$\lambda\left(\left(1-\alpha\right)p_r - c\right) + \delta\pi_r\left(\lambda\right). \tag{22}$$

Since $\pi_r\left(\lambda\right)$ is concave, this expression admits a unique maximum in λ given p_r.

Suppose that the equilibrium λ is larger than the one that maximizes the risky seller's profits (22) given p_r. Then the risky seller can simply restrict its p_r price offer to a lower number of period-1 buyers, thereby strictly increasing profits. Suppose instead that the equilibrium λ is smaller than the one that maximizes the risky seller's profits (22) given p_r. The risky seller can then very slightly decrease p_r: all period-1 buyers now strictly prefer the risky product over the safe product so are

[9] In particular, the prices p_r and p_s are offered to all buyers without any restrictions in equilibrium.

willing to buy from the risky seller. However, the risky seller can still restrict its price to be available only to the first $\lambda' > \lambda$ buyers that come to it. In this way, since the change in price is infinitesimal, the risky seller can strictly increase its profits. Thus, the candidate equilibrium λ must maximize the risky seller's total discounted profits given p_r, i.e.

$$(1-\alpha)\, p_r - c + \delta \pi_r' (\lambda) = 0. \tag{23}$$

Using a very similar reasoning, the equilibrium λ must also maximize the safe seller's total discounted profits $(1-\lambda)\left((1-\alpha)\, p_s - c\right) + \delta \pi_s (\lambda)$, which implies

$$-\left((1-\alpha)\, p_s - c\right) + \delta \pi_s' (\lambda) = 0. \tag{24}$$

Like for the risky seller, this also implies that equilibrium profits for the safe seller are positive.

Combining conditions (21), (23) and (24), the equilibrium λ and prices (p_r, p_s) must satisfy

$$\begin{aligned} \pi_s' (\lambda) + \pi_r' (\lambda) &= \left(\frac{1-\alpha}{\delta}\right) (u_s - \overline{u}_r) \\ p_r &= \frac{c - \delta \pi_r' (\lambda)}{1-\alpha} \\ p_s &= \frac{c + \delta \pi_s' (\lambda)}{1-\alpha}. \end{aligned}$$

Online Appendix

Andrei Hagiu[*] and Julian Wright[†]

In this online appendix, we provide proofs for additional results and claims referred to in the paper "Platforms and the exploration of new products" (hereafter, "the main paper").

A Private signals

In this online appendix we consider what happens when x is a buyer-specific shock (rather than a public signal as in the main paper) and sellers cannot charge prices conditional on x (either because x is privately observed by buyers or, even if it is publicly observed, sellers cannot price discriminate). In this case, the truth-or-noise model captures that the more buyers explore in period 1, the higher the chance that the signal received by any given buyer in period 2 is based on a matching period-1 buyer that shares the same preferences (and so reveals the true value for that period-2 buyer). Our analysis here also applies to the case when x is the same for all buyers, but is not observed by sellers.

First, note that the case with competing risky and competing safe sellers remains unchanged, given sellers do not extract any of the buyer' surplus anyway. There is no exploration in equilibrium and the platform is indifferent to its level.

For the other cases, the only thing that changes is the period-2 expected profit for the seller(s) with market power. Instead of extracting the entire upside from exploration, the seller(s) with market power only capture a fraction of it, reflecting that demand is now downward sloping.

Consider first the case with a risky seller and competing safe sellers. In period 2, a buyer can purchase from the safe sellers at $\frac{c}{1-\alpha}$ and obtain $u_s - \frac{c}{1-\alpha}$. If the buyer that receives the signal x buys from the risky seller, the buyer expects to obtain $E_\lambda [u_r|x] - p_r^2$, where p_r^2 represents the risky seller's period 2 price. Thus, the buyer with signal x will purchase from the risky seller if and only if $x \geq \overline{u}_r + \frac{u_s - \overline{u}_r + p_r^2 - \frac{c}{1-\alpha}}{F(\lambda)}$. The risky seller's period-2 expected profit is then

$$\widetilde{\pi}_r(\lambda) = (1-\alpha) \max_{p_r^2} \left\{ \left(p_r^2 - \frac{c}{1-\alpha} \right) \left(1 - G\left(\overline{u}_r + \frac{u_s - \overline{u}_r + p_r^2 - \frac{c}{1-\alpha}}{F(\lambda)} \right) \right) \right\}.$$

Assumption (2) in the main paper implies that this profit is positive for any λ. Moreover, $\widetilde{\pi}_r(\lambda)$ is increasing in λ since a higher $F(\lambda)$ increases demand for any price $p_r^2 > \frac{c}{1-\alpha}$. Denote by $p_r^{2*}(\lambda)$ the price chosen by the risky seller in period 2.

[*]Boston University Questrom School of Business
[†]Department of Economics, National University of Singapore

The platform's period-2 expected profit is

$$\begin{aligned}
&\alpha\left(\frac{c}{1-\alpha}G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+p_r^{2*}(\lambda)-\frac{c}{1-\alpha}}{F(\lambda)}\right)+p_r^{2*}(\lambda)\left(1-G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+p_r^{2*}(\lambda)-\frac{c}{1-\alpha}}{F(\lambda)}\right)\right)\right)\\
=&\ \alpha\left(\left(p_r^{2*}(\lambda)-\frac{c}{1-\alpha}\right)\left(1-G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+p_r^{2*}(\lambda)-\frac{c}{1-\alpha}}{F(\lambda)}\right)\right)+\frac{c}{1-\alpha}\right)\\
=&\ \frac{\alpha}{1-\alpha}\left(\widetilde{\pi}_r(\lambda)+c\right).
\end{aligned}$$

Since the period-1 profits for the seller and the platform are identical to those in Section 3.2 in the main paper, $\widetilde{\pi}_r(\lambda)$ is increasing in λ, and the relationship between the seller's total expected profit and the platform's expected profit is the same as in Section 3.2 in the main paper, Proposition 2 in the main paper continues to hold.

Next, consider the case with a safe seller and competing risky sellers. The safe seller's period-2 expected profit is

$$\widetilde{\pi}_s(\lambda)=(1-\alpha)\max_{p_s^2}\left\{\left(p_s^2-\frac{c}{1-\alpha}\right)G\left(\overline{u}_r+\frac{u_s-\overline{u}_r-p_s^2+\frac{c}{1-\alpha}}{F(\lambda)}\right)\right\}.$$

This profit is clearly positive for any λ. Moreover, $\widetilde{\pi}_s(\lambda)$ is increasing in λ. And the platform's period-2 expected profit is

$$\frac{\alpha}{1-\alpha}\left(\widetilde{\pi}_s(\lambda)+c\right).$$

Since the period-1 profits for the seller and the platform are identical to those in Section 3.3 in the main paper, $\widetilde{\pi}_r(\lambda)$ is increasing in λ, and the relationship between the seller's total expected profit and the platform's expected profit is the same as in Section 3.3 in the main paper, Proposition 3 in the main paper continues to hold.

Finally, consider the case with one risky and one safe seller. Denote by $p_r^*(\lambda)$ and $p_s^*(\lambda)$ the risky and safe seller's period-2 equilibrium prices. They solve:

$$\begin{aligned}
p_r^{2*}(\lambda) &= \arg\max_{p_r^2}\left\{\left(p_r^2-\frac{c}{1-\alpha}\right)\left(1-G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+p_r^2-p_s^{2*}(\lambda)}{F(\lambda)}\right)\right)\right\}\\
p_s^{2*}(\lambda) &= \arg\max_{p_s^2}\left\{\left(p_s^2-\frac{c}{1-\alpha}\right)G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+p_r^{2*}(\lambda)-p_s^2}{F(\lambda)}\right)\right\}.
\end{aligned}$$

The two sellers' period-2 expected profits are then

$$\begin{aligned}
\widetilde{\pi}_r(\lambda) &= (1-\alpha)\left(p_r^{2*}(\lambda)-\frac{c}{1-\alpha}\right)\left(1-G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+p_r^{2*}(\lambda)-p_s^{2*}(\lambda)}{F(\lambda)}\right)\right)\\
\widetilde{\pi}_s(\lambda) &= (1-\alpha)\left(p_s^{2*}(\lambda)-\frac{c}{1-\alpha}\right)G\left(\overline{u}_r+\frac{u_s-\overline{u}_r+p_r^{2*}(\lambda)-p_s^{2*}(\lambda)}{F(\lambda)}\right).
\end{aligned}$$

Note that an increase in λ makes both sellers' demands less sensitive to the difference in their prices, which should naturally result in higher equilibrium prices and profits (this is the same property as in the main paper). Here, we assume $G(.)$ is sufficiently well-behaved such that this property holds (e.g. it holds when $G(.)$ is log-concave).

The period-1 pricing game is then the same as in Section 3.4 in the main paper, so Lemma 1 in the main paper continues to apply, after replacing $\pi_r(\lambda)$ and $\pi_s(\lambda)$ by $\widetilde{\pi}_r(\lambda)$ and $\widetilde{\pi}_s(\lambda)$. Furthermore, the platform's period-2 expected profit is now

$$\begin{aligned} & \alpha \left(p_r^{2*}(\lambda) \left(1 - G\left(\overline{u}_r + \frac{u_s - \overline{u}_r + p_r^{2*}(\lambda) - p_s^{2*}(\lambda)}{F(\lambda)} \right) \right) + p_s^{2*}(\lambda) G\left(\overline{u}_r + \frac{u_s - \overline{u}_r + p_r^{2*}(\lambda) - p_s^{2*}(\lambda)}{F(\lambda)} \right) \right) \\ = & \frac{\alpha}{1-\alpha} \left(\widetilde{\pi}_r(\lambda) + \widetilde{\pi}_s(\lambda) + c \right), \end{aligned}$$

so the platform's profit expression (13) in the main paper also continues to apply, after replacing $\pi_r(\lambda)$ and $\pi_s(\lambda)$ by $\widetilde{\pi}_r(\lambda)$ and $\widetilde{\pi}_s(\lambda)$. Consequently, Proposition 4 in the main paper also continues to hold.

B Explicit platform steering

We consider a specific steering technology that the platform can use to steer buyers in period 1. Suppose prior to sellers setting prices for period-1 buyers, the platform can make a costly investment to temporary increase the perceived utility of buying a unit of the risky product vs. a unit of the safe product. Specifically, the platform can change the utility of buying the risky product in period 1 to $\overline{u}_r + \frac{\Delta}{2}$ and the utility of buying the safe product to $u_s - \frac{\Delta}{2}$ by incurring the cost $C(\Delta)$, where $C(0) = 0$, $C(\Delta)$ is increasing for $\Delta > 0$ and $C(\Delta)$ is decreasing for $\Delta < 0$ (i.e. steering in either direction is costly). With this specification, a positive Δ represents the platform steering towards the risky product and a negative Δ represents the platform steering towards the safe product. Then the game unfolds as before.

For a given λ, the second-period analysis is unchanged. Thus, we can focus on the first-period analysis for all four seller market structures. The first-period analysis will depend on two effects. One is the direct effect of Δ on the platform's period 1 payoffs, holding constant λ. The direct effect consists of the cost of steering C as well as any change in period-1 prices due to the change in Δ (e.g. if the utility of the risky product is increased, then the price of the risky product can be set higher in period 1). The second effect is the strategic effect of Δ through λ. This captures the effect on the platform's profit through any change in the level of exploration, holding constant period 1 prices. In the analysis throughout the paper, we have focused on this strategic effect when determining whether the platform prefers λ to be higher or lower than the equilibrium level λ^*. To work out whether the platform will steer and in what direction with this steering technology, we need to combine these two effects. To avoid that our steering results are driven by the direct effect, we assume the cost of steering is sufficiently high that the platform will never find it profitable to steer (in either direction) unless there is some gain to changing the level of exploration (i.e. the strategic effect is non-positive).

We consider the four different benchmark settings covered in Sections 3.1-3.4.

1. *Competitive sellers of both types.* Due to Bertrand-style price competition in both periods,

prices are set to recover cost, so there is no strategic effect of steering on the platform's profits. Meanwhile, with prices set to recover costs, the direct effect of steering is just the cost of steering, which is positive. Thus, the platform will not steer any buyers and there will be no exploration.

2. *Single seller of risky product.* Following the same argument as in Section 3.2, but modifying the utilities in period 1, the risky seller's total expected profit is

$$-\lambda(1-\alpha)(u_s - \overline{u}_r - \Delta) + \delta\pi_r(\lambda) \tag{B.1}$$

and the platform's expected profit is

$$\frac{\alpha}{1-\alpha}\left((1+\delta)c - \lambda(1-\alpha)(u_s - \overline{u}_r - \Delta) + \delta\pi_r(\lambda)\right) - C(\Delta). \tag{B.2}$$

Note (B.1) together with the property that $\pi_r''(\lambda) < 0$ implies λ is increasing in Δ. Since λ is chosen by the risky seller in a way that maximizes (B.2), there is no strategic reason to change Δ to either increase or decrease λ.[1] The lack of a strategic effect of steering just captures the result in Proposition 2, that the platform prefers no change in exploration relative to the (positive) equilibrium level holding first-period prices as given. The direct effect of Δ on the platform's profit is $\alpha\lambda\Delta - C(\Delta)$, which is negative for all $\Delta < 0$. For $\Delta > 0$, the direct effect of increasing Δ by one unit is just $\alpha\lambda - C'(\Delta)$. A higher Δ temporarily shifts more perceived surplus to the risky seller's product, which reduces the amount the risky seller has to lower its price in period 1 to attract buyers. This benefits the platform since period 1 revenue is higher, and so its share of revenue is higher. But given our assumption that this direct effect is unprofitable, then steering in either direction will not be profitable. A sufficient condition to ensure this is that $|C'(\Delta)| > \alpha$ for any $\Delta \neq 0$.

3. *Single seller of safe product.* A parallel argument to case 2 applies and we conclude that steering in either direction remains unprofitable.

4. *Single seller of each type of product.* Following the analysis in Section 3.4, the equilibrium level of exploration λ^* induced by Δ is determined by

$$\pi_s'(\lambda^*) + \pi_r'(\lambda^*) = \frac{1-\alpha}{\delta}(u_s - \overline{u}_r - \Delta).$$

This implies λ^* is increasing in Δ. Meanwhile, the platform's objective function is

$$\alpha\left(\lambda^* p_r^* + (1-\lambda^*)p_s^* + \frac{\delta}{1-\alpha}(\pi_r(\lambda^*) + \pi_s(\lambda^*) + c)\right) - C(\Delta),$$

[1] The platform also cannot do better by setting $\Delta \geq u_s - \overline{u}_r$ to induce $\lambda = 1$. We know that increasing Δ to move from λ^* to $\lambda = 1$ lowers the risky seller's profit (by our assumption that $F'(1)$ is sufficiently small so that the equilibrium level of λ is interior), so this also lowers the platform's profit given it is proportional to the risky seller's profit, as shown above.

where (from Lemma 1)

$$p_r^* = \frac{c - \delta\pi_r'(\lambda^*)}{1-\alpha} \text{ and } p_s^* = \frac{c + \delta\pi_s'(\lambda^*)}{1-\alpha}.$$

This implies

$$\begin{aligned}
\frac{dp_r^*}{d\Delta} &= -\frac{\delta\pi_r''(\lambda^*)}{1-\alpha}\frac{d\lambda^*}{d\Delta} = \frac{\pi_r''(\lambda^*)}{\pi_s''(\lambda^*) + \pi_r''(\lambda^*)} \\
\frac{dp_s^*}{d\Delta} &= \frac{\delta\pi_s''(\lambda^*)}{1-\alpha}\frac{d\lambda^*}{d\Delta} = -\frac{\pi_s''(\lambda^*)}{\pi_s''(\lambda^*) + \pi_r''(\lambda^*)}.
\end{aligned}$$

Thus, the derivative of the platform's profit with respect to Δ can be written

$$\begin{aligned}
& \alpha\left(p_r^* - p_s^* + \frac{\delta}{1-\alpha}\left(\pi_r'(\lambda^*) + \pi_s'(\lambda^*)\right)\right)\frac{d\lambda^*}{d\Delta} + \alpha\left(\lambda^*\frac{dp_r^*}{d\Delta} + (1-\lambda^*)\frac{\Delta p_s^*}{d\Delta}\right) - C'(\Delta) \\
= \; & \alpha\left(\lambda^*\frac{dp_r^*}{d\Delta} + (1-\lambda^*)\frac{\Delta p_s^*}{d\Delta}\right) - C'(\Delta) \\
= \; & \alpha\left(\lambda^* - \frac{\pi_s''(\lambda^*)}{\pi_s''(\lambda^*) + \pi_r''(\lambda^*)}\right) - C'(\Delta)
\end{aligned}$$

Thus, provided $|C'(\Delta)| > \alpha$ and since $0 < \frac{\pi_s''(\lambda^*)}{\pi_s''(\lambda^*)+\pi_r''(\lambda^*)} < 1$, the platform's profit is lower if $\Delta < 0$ or $\Delta > 0$ compared to $\Delta = 0$. Therefore the platform does not want to steer buyers in either direction.

C Uniform pricing and costly steering by sellers

Suppose there is one risky seller and identical safe sellers, and the risky seller can only set a uniform price. We want to show that we obtain the same result as in Section 3.2 if we allow the risky seller to incur a cost for each buyer that it attracts to buy its product, where the cost to attract the buyer is increasing in the net surplus shortfall of the risky seller's product.

Specifically, suppose the risky seller can invest $K(\Delta)$ in order to get a period-1 buyer to buy the risky product, where $K(\Delta)$ is increasing and continuously differentiable for $\Delta \geq 0$, and $\Delta = (u_s - p_s) - (\overline{u}_r - p_r)$ is the difference in net utility between the two products in period 1. We assume $K(\Delta) = 0$ for all $\Delta \leq 0$ so there is no cost to steer buyers to the risky product if the net utility offered by the two choices is identical or the risky seller offers higher net utility in period 1. Consistent with this specification, we no longer assume that when buyers are indifferent between the two products, buyers all buy the product from the seller with market power, but rather that they can be steered by the risky seller to the seller's desired level of exploration at no cost.[2] Our formulation represents the limit case of this analysis taking p_r to the limit where buyers are indifferent between the two

[2] In practice, to achieve this, the risky seller can set a slightly higher p_r than the price that makes buyers indifferent (so that buyers strictly prefer to buy from the safe sellers), but then steer λ buyers to buy from it at a very small cost $K(\Delta)$ given that Δ is very close to zero.

products. We also assume that $K'(\Delta) > 1-\alpha$ for all $\Delta \geq 0$, so that if the surplus difference increases by one unit, the cost of steering the buyer will increase by more than $1-\alpha$, where recall $\alpha > 0$ is the percentage commission extracted by the platform. Obviously, this assumption means that discounting the price to a set of buyers is a more efficient way to get these buyers to explore than steering them through the costly investment. The timing and payoff specification of the modified game we consider is as follows:

Period 1a Sellers set (uniform) prices for period-1 buyers.

Period 1b The risky seller chooses how many buyers to steer to choose its product. If it wants to steer λ buyers, it has to incur the cost $\lambda K(\Delta)$.

Period 1c Then λ buyers purchase from the risky seller and the remaining $1-\lambda$ buyers purchase from the safe seller.

Period 2a The signal x is realized and observed by all players.

Period 2b Sellers set prices for period-2 buyers.

Period 2c Period-2 buyers make their purchase decisions.

Given Bertrand competition, the safe sellers always price at $\frac{c}{1-\alpha}$ in both periods. Given λ buyers explore in period 1, the period 2 analysis is identical to the benchmark case.

In period 1, the risky seller considers whether to induce λ buyers to try its product. The cost to the risky seller of inducing λ buyers to explore through direct steering is $\lambda K\left(\left(u_s - \frac{c}{1-\alpha}\right) - (\overline{u}_r - p_r)\right)$. Provided $p_r \geq \frac{c}{1-\alpha} + \overline{u}_r - u_s$, only buyers that are induced to explore will purchase at this price. This means that to induce λ buyers to explore, the risky seller has a choice between lowering its price p_r for the exploring buyers all the way to $\frac{c}{1-\alpha} + \overline{u}_r - u_s$ (which eliminates the need for incurring direct steering costs), or keeping a larger p_r and incurring positive steering costs. The first-period profit reflecting this choice for a given λ is therefore

$$\lambda\left((1-\alpha)\,p_r - c - K\left(\left(u_s - \frac{c}{1-\alpha}\right) - (\overline{u}_r - p_r)\right)\right),$$

where $p_r \geq \frac{c}{1-\alpha} + \overline{u}_r - u_s$. Given our assumption that $K'(\Delta) > 1-\alpha$ for all $\Delta \geq 0$, the risky seller will always do best setting $p_r = \frac{c}{1-\alpha} + \overline{u}_r - u_s$ and not incurring any direct steering costs to get the λ buyers to buy from it.

Thus, the problem becomes identical to the one in the baseline setting in which the risky seller sets the lower price of $\frac{c}{1-\alpha} + \overline{u}_r - u_s$ for the λ buyers which it sells to, and the rest of the analysis applies as in Section 3.2. This same logic also applies to other two cases in which at least one type of seller has market power.

D Market expansion through elastic buyer participation

Suppose the total number of buyers who participate in period 2 depends positively on their expected surplus (this can be obtained by assuming buyers incur a heterogeneous participation cost). Specifically, assume the total number of buyers in period 2 is $N\left(U_2\left(\lambda\right)\right)$, where $N\left(.\right)$ is an increasing function and $U_2\left(\lambda\right)$ is the period-2 buyer expected surplus as a function of the level of exploration λ. This presumes that period-2 buyers observe λ but not x prior to joining the platform. For example, buyers may be able to see how many reviews there are before incurring the cost to use the platform, but only observe the signal x after reading the details of the reviews (i.e. incurring the cost to use the platform). The function $U_2\left(\lambda\right)$ depends on the seller market structure and was determined in Section 4.1.

With one risky seller and competing safe sellers, we have $U_2\left(\lambda\right) = u_s - \frac{c}{1-\alpha}$. This means $N\left(U_2\left(\lambda\right)\right)$ does not depend on λ, so our benchmark result remains unchanged: the platform is content with the equilibrium level of exploration. With one safe seller and competing risky sellers, we have $U_2\left(\lambda\right) = E_x\left[E_\lambda\left[u_r|x\right] - \frac{c}{1-\alpha}\right] = \overline{u}_r - \frac{c}{1-\alpha}$. This means $N\left(U_2\left(\lambda\right)\right)$ does not depend on λ, so once again our benchmark result still holds: the platform is content with the equilibrium level of exploration.

With competing sellers of both types, we have

$$U_2\left(\lambda\right) = u_s G\left(\overline{u}_r + \frac{u_s - \overline{u}_r}{F\left(\lambda\right)}\right) + \int_{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}}^{u_H} E_\lambda\left[u_r|x\right] dG\left(x\right) - \frac{c}{1-\alpha},$$

which is increasing in λ. The platform's objective function is $\alpha\left(\frac{c}{1-\alpha} + \delta N\left(U_2\left(\lambda\right)\right)\frac{c}{1-\alpha}\right)$, so is increasing in λ. This means the platform would like to see maximum exploration, whereas the equilibrium level is zero. Thus, in this case there is insufficient exploration, just like in Section 4.1.

With one seller of each type, we have

$$U_2\left(\lambda\right) = \int_{u_L}^{\overline{u}_r + \frac{u_s - \overline{u}_r}{F(\lambda)}} E_\lambda\left[u_r|x\right] dG\left(x\right) + u_s\left(1 - G\left(\overline{u}_r + \frac{u_s - \overline{u}_r}{F\left(\lambda\right)}\right)\right) - \frac{c}{1-\alpha},$$

which is decreasing in λ.

Lemma 1 is easily extended to this case (the only difference is that second period profits are $N\left(U_2\left(\lambda\right)\right)\pi_s\left(\lambda\right)$ and $N\left(U_2\left(\lambda\right)\right)\pi_r\left(\lambda\right)$), so the equilibrium level of exploration λ^* is given by

$$\left.\frac{d\left(N\left(U_2\left(\lambda\right)\right)\left(\pi_s\left(\lambda\right) + \pi_r\left(\lambda\right)\right)\right)}{d\lambda}\right|_{\lambda=\lambda^*} = \frac{1-\alpha}{\delta}\left(u_s - \overline{u}_r\right)$$

The platform's objective function (taking as given the first period equilibrium prices p_r^* and p_s^*) is

$$\alpha\left(\lambda p_r^* + \left(1-\lambda\right)p_s^* + \frac{\delta}{1-\alpha}N\left(U_2\left(\lambda\right)\right)\left(\pi_r\left(\lambda\right) + \pi_s\left(\lambda\right) + c\right)\right).$$

Taking into account that in equilibrium we must have $p_r^* - p_s^* = \overline{u}_r - u_s$ and recalling that $N\left(U_2\left(\lambda\right)\right)$ is decreasing in this case, it is straightforward to see that the platform prefers a lower level of expectation than that prevailing in equilibrium, just like in Section 4.1.

In summary, the results in (i)-(iii) in Proposition 5 continue to apply here.

E Proof of Proposition 6

The case with competing sellers of both types and the one with a single risky seller and competing safe sellers were analyzed in the main text. Here, we provide the analysis for the other two cases.

Consider first the case with a single safe seller facing competitive risky sellers. Paralleling the analysis in Section 3.3 in the main paper, the safe seller's period-2 expected profit is $\pi_s(\lambda, N(\lambda))$, where

$$\begin{aligned}\pi_s(\lambda, N) &\equiv (1-\alpha)(1+N)\int_{u_L}^{\overline{u}_r+\frac{u_s-\overline{u}_r}{F(\lambda)}} (u_s - E_\lambda[u_r|x])\,dG(x) \\ &= (1-\alpha)(1+N)(u_s-\overline{u}_r) + \pi_r(\lambda, N).\end{aligned} \tag{E.1}$$

Thus,

$$\frac{\partial \pi_s(\lambda, N)}{\partial \lambda} = \frac{\partial \pi_r(\lambda, N)}{\partial \lambda} \tag{E.2}$$

$$\frac{\partial \pi_s(\lambda, N)}{\partial N} = \frac{\partial \pi_r(\lambda, N)}{\partial N} + (1-\alpha)(u_s-\overline{u}_r), \tag{E.3}$$

so both derivatives are positive. The safe seller's total expected profit is then

$$(1-\alpha)(1-\lambda)(u_s-\overline{u}_r) + \delta\pi_s(\lambda, N(\lambda)). \tag{E.4}$$

Taking the total derivative of (E.4) with respect to λ we obtain that the equilibrium level of exploration induced by the safe seller is the solution to

$$\frac{\partial \pi_s(\lambda, N(\lambda))}{\partial \lambda} + \frac{\partial \pi_s(\lambda, N(\lambda))}{\partial N} N'(\lambda) - \frac{1-\alpha}{\delta}(u_s-\overline{u}_r) = 0. \tag{E.5}$$

Meanwhile, the platform's expected profit can be written as

$$\alpha\left((1-\lambda)(u_s-\overline{u}_r) + \frac{c}{1-\alpha}\right) + \frac{\delta\alpha}{1-\alpha}(\pi_s(\lambda, N) + (1+N)c) - C(N). \tag{E.6}$$

Given λ, the platform chooses N to maximize (E.6), implying $N(\lambda)$ is the solution to

$$\frac{\alpha}{1-\alpha}\left(\frac{\partial \pi_s(\lambda, N)}{\partial N} + c\right) = \frac{C'(N)}{\delta}.$$

From (E.3), $\frac{\partial^2 \pi_s(\lambda,N)}{\partial N \partial \lambda} = \frac{\partial^2 \pi_r(\lambda,N)}{\partial N \partial \lambda}$, implying that $N(\lambda)$ is the same as the one determined by (17) in the main paper. Moreover, using the envelope theorem, the derivative of the platform's profit with respect to λ (taking as given the sellers' period 1 equilibrium prices) is proportional to

$$\frac{\partial \pi_s (\lambda, N(\lambda))}{\partial \lambda} - \frac{1-\alpha}{\delta} (u_s - \overline{u}_r)$$

Comparing this expression with the left-hand side of (E.5), the difference is the term $\frac{\partial \pi_s(\lambda, N(\lambda))}{\partial N} N'(\lambda)$, which is strictly positive given $\frac{\partial \pi_s(\lambda, N(\lambda))}{\partial N} > 0$ and $N'(\lambda) > 0$. This implies the platform prefers less exploration than the equilibrium level.

Now consider the case with a single safe seller and a single risky seller. In period 2, the risky seller's expected profit is $\pi_r(\lambda, N(\lambda))$, while the safe seller's expected profit is $\pi_s(\lambda, N(\lambda))$, where $\pi_r(\lambda, N)$ and $\pi_s(\lambda, N)$ are given by (14) in the main paper and (E.1), and the function $N(\lambda)$ is determined by the platform.

Taking as given the sellers' period-1 equilibrium prices (p_r^*, p_s^*) (which we will determine below), the platform's expected profit is

$$\alpha \left(\lambda p_r^* + (1-\lambda) p_s^* + \frac{\delta}{1-\alpha} (\pi_s(\lambda, N) + \pi_r(\lambda, N) + (1+N) c) \right) - C(N). \tag{E.7}$$

Taking the first-order condition in N for a given λ, we obtain that $N(\lambda)$ is defined by

$$\frac{\alpha}{1-\alpha} \left(\frac{\partial \pi_r(\lambda, N)}{\partial N} + \frac{\partial \pi_s(\lambda, N)}{\partial N} + c \right) = \frac{C'(N)}{\delta}.$$

Since $\frac{\partial^2 \pi_r(\lambda, N)}{\partial N \partial \lambda} = \frac{\partial^2 \pi_s(\lambda, N)}{\partial N \partial \lambda} > 0$, $N(\lambda)$ is increasing. And since $\pi_r(\lambda, N)$ and $\pi_s(\lambda, N)$ are both increasing in λ, we can conclude that $\pi_r(\lambda, N(\lambda))$ and $\pi_s(\lambda, N(\lambda))$ are increasing in λ. We will also assume second order conditions hold, such that the functions $\pi_r(\lambda, N(\lambda))$ and $\pi_s(\lambda, N(\lambda))$ are concave in λ.

We can then obtain a very similar result to Lemma 1 in the main paper.

Lemma *With market expansion and a single seller of each type, the equilibrium level of exploration* λ^* *is defined by*

$$\left.\frac{d\pi_s(\lambda, N(\lambda))}{d\lambda}\right|_{\lambda=\lambda^*} + \left.\frac{d\pi_r(\lambda, N(\lambda))}{d\lambda}\right|_{\lambda=\lambda^*} = \frac{1-\alpha}{\delta} (u_s - \overline{u}_r). \tag{E.8}$$

We now wish to determine whether the platform prefers more or less exploration relative to this equilibrium level, taking as given the prices chosen by sellers. From (E.7), using the envelope theorem and $p_s^* - p_r^* = u_s - \overline{u}_r$, the derivative of the platform's profit in λ evaluated at $N = N(\lambda)$ is proportional to

$$\left.\frac{\partial \pi_s(\lambda, N)}{\partial \lambda}\right|_{N=N(\lambda)} + \left.\frac{\partial \pi_r(\lambda, N)}{\partial \lambda}\right|_{N=N(\lambda)} - \frac{1-\alpha}{\delta} (u_s - \overline{u}_r). \tag{E.9}$$

Comparing (E.9) with (E.8) and recalling that $\pi_r(\lambda, N)$ and $\pi_s(\lambda, N)$ are increasing in N and $N'(\lambda) > 0$, we can conclude that the platform prefers less exploration relative to the equilibrium level.

F Proof of Proposition 7

The case with competing sellers of each type and the case with one risky seller facing competing safe sellers were treated in the main paper. Here we provide the analysis for the remaining two cases.

Consider first the case with one safe seller and competing risky sellers. The risky sellers price at $\frac{c_r}{1-\alpha_r}$ in both periods. For the safe seller to make any sales in period 1, the maximum price it can charge is

$$p_s = u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r}.$$

In period 2, the safe seller's profit is

$$\pi_s(\lambda) \equiv (1-\alpha_s) \int_{u_L}^{\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)}} \left(u_s - E_\lambda[u_r|x] + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s} \right) dG(x). \tag{F.1}$$

Thus, the safe seller's total expected profit as a function of the level of exploration chosen is

$$\Pi_s(\lambda) \equiv (1-\lambda)(1-\alpha_s)\left(u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s} \right) + \delta \pi_s(\lambda).$$

The level of exploration chosen by the safe seller is determined by the first-order condition $\Pi_s'(\lambda) = 0$.

Meanwhile, taking as given the sellers' period-1 equilibrium prices, the platform's expected profit is

$$\begin{aligned}
&\alpha_r \lambda \frac{c_r}{1-\alpha_r} + \alpha_s (1-\lambda)\left(u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} \right) \\
&+\delta \left(\alpha_s \int_{u_L}^{\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)}} \left(u_s - E_\lambda[u_r|x] + \frac{c_r}{1-\alpha_r} \right) dG(x) + \alpha_r \left(1 - G\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)} \right) \right) \right)
\end{aligned}$$

which can be re-written as

$$\frac{\alpha_s}{1-\alpha_s}\Pi_s(\lambda) + \left(\frac{\alpha_r c_r}{1-\alpha_r} - \frac{\alpha_s c_s}{1-\alpha_s} \right)\left(\lambda + \delta\left(1 - G\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)} \right) \right) \right) + \frac{\alpha_s c_s}{1-\alpha_s} + \delta \frac{\alpha_s c_s}{1-\alpha_s}.$$

Under assumption (19) in the main paper, the function $\lambda + \delta\left(1 - G\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)} \right) \right)$ is increasing in λ, so the platform prefers more exploration relative to the equilibrium level obtained by maximizing the safe seller's total expected profit $\Pi_s(\lambda)$ if and only if $\frac{\alpha_r c_r}{1-\alpha_r} > \frac{\alpha_s c_s}{1-\alpha_s}$.

Next, consider the case with one seller of each type. Using the second period profit expressions $\pi_s(\lambda)$ and $\pi_r(\lambda)$ defined by (20) in the main paper and (F.1), we can write the equilibrium conditions that must be satisfied by λ^* and by the sellers' period-1 prices p_r^* and p_s^*:

$$\begin{aligned}
(1-\alpha_r)p_r^* - c_r + \delta \pi_r'(\lambda) &= 0 \\
-(1-\alpha_s)p_s^* + c_s + \delta \pi_s'(\lambda) &= 0 \\
p_r^* - p_s^* &= \overline{u}_r - u_s.
\end{aligned}$$

Combining these equations to eliminate the prices p_r^* and p_s^*, we obtain that the equilibrium level of exploration must satisfy

$$\frac{\pi_r'(\lambda^*)}{1-\alpha_r} + \frac{\pi_s'(\lambda^*)}{1-\alpha_s} = \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{\delta}. \tag{F.2}$$

We can also combine the equations above to obtain

$$\alpha_r p_r^* - \alpha_s p_s^* = \delta\left(\pi_r'(\lambda^*) + \pi_s'(\lambda^*)\right) + (\overline{u}_r - u_s) + (c_s - c_r). \tag{F.3}$$

Taking as given the sellers' period 1 equilibrium prices, the platform's expected profit is

$$\alpha_r \lambda p_r^* + \alpha_s (1-\lambda_s) p_s^* + \delta \begin{pmatrix} \alpha_r \int_{\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)}}^{u_H} \left(E_\lambda[u_r|x] - u_s + \frac{c_s}{1-\alpha_s}\right) dG(x) \\ +\alpha_s \int_{u_L}^{\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)}} \left(u_s - E_\lambda[u_r|x] + \frac{c_r}{1-\alpha_r}\right) dG(x) \end{pmatrix}$$

$$= \alpha_r \lambda p_r^* + \alpha_s (1-\lambda) p_s^* + \delta \begin{pmatrix} \frac{\alpha_r}{1-\alpha_r}\pi_r(\lambda) + \frac{\alpha_r c_r}{1-\alpha_r}\left(1 - G\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)}\right)\right) \\ +\frac{\alpha_s}{1-\alpha_s}\pi_s(\lambda) + \frac{\alpha_s c_s}{1-\alpha_s} G\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)}\right) \end{pmatrix}.$$

Taking the derivative with respect to λ of the platform's profit and evaluating it at the equilibrium level of exploration $\lambda = \lambda^*$, we obtain

$$\alpha_r p_r^* - \alpha_s p_s^* + \delta \begin{pmatrix} \frac{\alpha_r}{1-\alpha_r}\pi_r'(\lambda^*) + \frac{\alpha_s}{1-\alpha_s}\pi_s'(\lambda^*) \\ +\left(\frac{\alpha_r c_r}{1-\alpha_r} - \frac{\alpha_s c_s}{1-\alpha_s}\right) F'(\lambda^*)\left(\frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda^*)^2}\right) g\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda^*)}\right) \end{pmatrix}.$$

Using (F.2) and (F.3), this expression is proportional to

$$\begin{aligned} &\left(\pi_r'(\lambda^*) + \pi_s'(\lambda^*)\right) + \frac{(\overline{u}_r - u_s) + (c_s - c_r)}{\delta} + \left(\frac{\alpha_r}{1-\alpha_r}\pi_r'(\lambda^*) + \frac{\alpha_s}{1-\alpha_s}\pi_s'(\lambda^*)\right) \\ &+ \left(\frac{\alpha_r c_r}{1-\alpha_r} - \frac{\alpha_s c_s}{1-\alpha_s}\right) F'(\lambda^*)\left(\frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda^*)^2}\right) g\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda^*)}\right) \end{aligned}$$

$$\begin{aligned} = \; &\frac{(\overline{u}_r - u_s) + (c_s - c_r)}{\delta} + \frac{\pi_r'(\lambda^*)}{1-\alpha_r} + \frac{\pi_s'(\lambda^*)}{1-\alpha_s} \\ &+ \left(\frac{\alpha_r c_r}{1-\alpha_r} - \frac{\alpha_s c_s}{1-\alpha_s}\right) F'(\lambda^*)\left(\frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda^*)^2}\right) g\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda^*)}\right) \end{aligned}$$

$$= \left(\frac{\alpha_r c_r}{1-\alpha_r} - \frac{\alpha_s c_s}{1-\alpha_s}\right)\left(\frac{1}{\delta} + F'(\lambda^*)\left(\frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda^*)^2}\right) g\left(\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda^*)}\right)\right).$$

Thus, the derivative with respect to λ of the platform's profit evaluated at the equilibrium level of exploration is positive if and only if $\frac{\alpha_r c_r}{1-\alpha_r} > \frac{\alpha_s c_s}{1-\alpha_s}$. In other words, the platform prefers more exploration relative to the equilibrium level if $\frac{\alpha_r c_r}{1-\alpha_r} > \frac{\alpha_s c_s}{1-\alpha_s}$, and less exploration if $\frac{\alpha_r c_r}{1-\alpha_r} < \frac{\alpha_s c_s}{1-\alpha_s}$.

G Proof of Proposition 9

In period 2, if the signal is x, then the location y of the buyer indifferent between the two sellers is determined by

$$u_s - p_s - ty = E_\lambda\left[u_r|x\right] - p_r - t\left(1-y\right),$$

where

$$E_\lambda\left[u_r|x\right] \equiv F\left(\lambda\right)x + \left(1-F\left(\lambda\right)\right)\overline{u}_r.$$

Assuming the market is covered (i.e. u_s and $E_\lambda\left[u_r|u_L\right]$ are high enough relative to t), the resulting period 2 profits are

$$\begin{aligned}\pi_s\left(\lambda\right) &= \left(\left(1-\alpha\right)p_s - c\right)\left(\frac{1}{2} + \frac{u_s - E_\lambda\left[u_r|x\right] + p_r - p_s}{2t}\right)\\ \pi_r\left(\lambda\right) &= \left(\left(1-\alpha\right)p_r - c\right)\left(\frac{1}{2} + \frac{E_\lambda\left[u_r|x\right] - u_s + p_s - p_r}{2t}\right).\end{aligned}$$

Then equilibrium pricing in period 2 is

$$\begin{aligned}p_s^* &= \frac{c}{1-\alpha} + t + \frac{u_s - E_\lambda\left[u_r|x\right]}{3}\\ p_r^* &= \frac{c}{1-\alpha} + t + \frac{E_\lambda\left[u_r|x\right] - u_s}{3},\end{aligned}$$

resulting in period-2 equilibrium profits

$$\begin{aligned}\pi_s^*\left(\lambda|x\right) &= \left(1-\alpha\right)\frac{t}{2}\left(1 + \frac{u_s - E_\lambda\left[u_r|x\right]}{3t}\right)^2\\ \pi_r^*\left(\lambda|x\right) &= \left(1-\alpha\right)\frac{t}{2}\left(1 + \frac{E_\lambda\left[u_r|x\right] - u_s}{3t}\right)^2.\end{aligned}$$

From the perspective of period 1, we need to take expectation over x, so expected second-period profits are (assuming the cutoff number of consumers is interior for all x, i.e. that t is high enough, u_s and $E_\lambda\left[u_r|u_L\right]$ are not too different, and u_s and $E_\lambda\left[u_r|u_H\right]$ are not too different)

$$\begin{aligned}\pi_s^e\left(\lambda\right) &= \left(1-\alpha\right)\frac{t}{2}\int_{u_L}^{u_H}\left(1 + \frac{u_s - E_\lambda\left[u_r|x\right]}{3t}\right)^2 dG\left(x\right)\\ &= \left(1-\alpha\right)\frac{\left(3t + u_s - \overline{u}_r\right)^2 + F\left(\lambda\right)^2 Var\left(u_r\right)}{18t}\\ \pi_r^e\left(\lambda\right) &= \left(1-\alpha\right)\frac{t}{2}\int_{u_L}^{u_H}\left(1 - \frac{u_s - E_\lambda\left[u_r|x\right]}{3t}\right)^2 dG\left(x\right)\\ &= \left(1-\alpha\right)\frac{\left(3t + \overline{u}_r - u_s\right)^2 + F\left(\lambda\right)^2 Var\left(u_r\right)}{18t}.\end{aligned}$$

Note as in our benchmark model, more exploration increases the expected second-period profits of both the safe and the risky seller because $Var\left(u_r\right) > 0$.

Next consider period 1. With a slight abuse of notation, p_s and p_r now denote period 1 prices. Profits are then

$$\begin{aligned}\pi_s &= ((1-\alpha)\,p_s - c)\,(1-\lambda) + \delta\pi_s^e(\lambda)\\ \pi_r &= ((1-\alpha)\,p_r - c)\,\lambda + \delta\pi_r^e(\lambda)\end{aligned}$$

where

$$\lambda = \frac{1}{2} + \frac{\overline{u}_r - u_s + p_s - p_r}{2t}$$

is the fraction of consumers that buy from the risky seller in period 1 (again, assuming a covered market in period 1). Then the first-order conditions in period 1 are

$$\begin{aligned}\frac{1}{2} + \frac{u_s - \overline{u}_r + \frac{c}{1-\alpha} + p_r - 2p_s}{2t} + \delta\frac{F(\lambda)\,F'(\lambda)}{18t^2}Var(u_r) &= 0\\ \frac{1}{2} + \frac{\overline{u}_r - u_s + \frac{c}{1-\alpha} + p_s - 2p_r}{2t} - \delta\frac{F(\lambda)\,F'(\lambda)}{18t^2}Var(u_r) &= 0,\end{aligned}$$

which imply period 1 equilibrium prices are

$$\begin{aligned}p_s^* &= \frac{c}{1-\alpha} + t + \frac{u_s - \overline{u}_r}{3} + \delta\frac{F(\lambda)\,F'(\lambda)}{27t}Var(u_r)\\ p_r^* &= \frac{c}{1-\alpha} + t + \frac{\overline{u}_r - u_s}{3} - \delta\frac{F(\lambda)\,F'(\lambda)}{27t}Var(u_r).\end{aligned}$$

The price effect due to exploration means the safe seller increases its price and the risky seller decreases its price, which induces more buyers to try the risky seller in period 1. Note that, unlike our benchmark model, here sellers have no need to offer different prices to different buyer segments since buyers are anyway heterogeneous in their preferences. Thus, sellers compete in uniform prices.

The equilibrium level of exploration is then determined by

$$\lambda^* = \frac{1}{2} - \frac{u_s - \overline{u}_r}{6t} + \delta\frac{F(\lambda^*)\,F'(\lambda^*)}{27t^2}Var(u_r). \tag{G.1}$$

Note that λ^* can be higher or lower than $\frac{1}{2}$. Assuming t is sufficiently high, (G.1) admits a unique solution λ^* which must lie within $(0,1)$.

The extent to which the safe product is valued higher than the risky product in period 1 $(u_s - \overline{u}_r)$ makes fewer buyers want to try the new seller ($\lambda^* < \frac{1}{2}$), but taking into account the effect of exploration on period 2 profits leads to a larger λ^* as both sellers want to induce more exploration (this makes the safe seller price less aggressively and the risky seller to price more aggressively in period 1). The latter effect can be smaller or larger than the former. Note also that λ^* is decreasing in $u_s - \overline{u}_r$ and increasing in the variance $Var(u_r)$, but the effect of t (the degree of differentiation) on λ^* is ambiguous.

Compare now the solution of (G.1) to the level of exploration that the platform wants. The platform's expected profits at the existing period 1 equilibrium prices can be written

$$\alpha\left(\lambda^* p_r^* + (1-\lambda^*)\, p_s^*\right) + \frac{\alpha\delta}{1-\alpha}\left(c + \pi_s^e(\lambda^*) + \pi_r^e(\lambda^*)\right).$$

Thus, if the platform can exogenously increase λ^* by an infinitesimal amount without affecting first-period prices, but taking into account how this influences second-period prices, then the impact on its profit is measured by

$$\begin{aligned}
& \alpha\left(p_r^* - p_s^*\right) + \frac{\alpha\delta}{1-\alpha}\left(\frac{\partial \pi_s^e(\lambda^*)}{\partial\lambda} + \frac{\partial \pi_r^e(\lambda^*)}{\partial\lambda}\right) \\
= \;& \alpha\left(\frac{2}{3}\left(\overline{u}_r - u_s\right) - \delta\frac{2F(\lambda^*)\,F'(\lambda^*)}{27t}\left(\int_{u_L}^{u_H} x^2 dG(x) - \overline{u}_r^2\right)\right) \\
& + \alpha\delta\frac{2F(\lambda^*)\,F'(\lambda^*)}{9t}\left(\int_{u_L}^{u_H} x^2 dG(x) - \overline{u}_r^2\right) \\
= \;& 4\alpha t\left(\frac{\overline{u}_r - u_s}{6t} + \delta\frac{F(\lambda^*)\,F'(\lambda^*)}{27t^2}\left(\int_{u_L}^{u_H} x^2 dG(x) - \overline{u}_r^2\right)\right) \\
= \;& 4\alpha t\left(\lambda^* - \frac{1}{2}\right),
\end{aligned}$$

where we have used (G.1) to obtain the last equality. This result implies that when sellers induce more/less than half of buyers to explore in equilibrium, the platform will want even more/less exploration.