Platforms and the exploration of new products*

Andrei Hagiu[†] and Julian Wright[‡]

June 14th, 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, if anything, seem to steer consumers towards established products and sellers.

JEL classification: D23, D86, L14

Keywords: platform design, product experimentation, steering buyers

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 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.

*We thank Doh-shin Jeon, Luis Cabral, Yi Zhu, participants at 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, during which we conducted a large part of the research presented in this paper.

[†]MIT Sloan School of Management, Cambridge, MA 02142, E-mail: ahagiu@mit.edu

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

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 the 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. Moreover, 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, the platform seldom has any incentive to induce more exploration, and under various reasonable conditions, may actually want *less* exploration than that induced by sellers. This explanation 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: the platform also cares about buyer surplus in additional to seller revenues, new sellers differ from established sellers in their marginal costs and/or in the share of revenue extracted by the platform, the platform can invest in marketing to attract more buyers onto the platform, and the new sellers get additional benefits outside the platform from increased exploration on the platform (e.g. increases sales through other channels). We are thus able to provide some guidance to platforms on when they should steer buyers towards established or new products/sellers, or when there is no need 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. 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 compare its desired level of exploration (as opposed to a social planner's) with the equilibrium level induced by sellers. Second, we perform this comparison 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 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.

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

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

This requires $F(0) > 0$, which was assumed above. Assumption (2) means 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. The idea is that the signal obtained in period 2 may reveal some useful information even if there is no buyer exploration (e.g. this may represent learning via other channels). 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. This is a technical assumption, the main role of which is to help ensure interior solutions.

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

[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.

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. 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). Finally, case 3 may 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. With a slight abuse of language, from now on we will say "safe sellers" when referring to sellers of the safe product and "risky sellers" when referring to 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

[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.

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 4.2 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 and determine λ.

Period 1b Period-1 buyers make their purchase decisions.

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 (positive), $F'(1)$ is sufficiently close to zero, and F'' is sufficiently negative so that the following function

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

is concave in λ. This ensures that the respective first-order conditions considered throughout the paper uniquely characterize optimal 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 like to steer more or less buyers to explore the risky product compared to the equilibrium outcome without any platform intervention. To determine this we start by characterizing the equilibrium level of λ (i.e., exploration) induced by sellers. We then examine whether the platform prefers a higher or a lower level of λ taking as given the sellers' prices in period 1 (since they are set at the same time as λ), but allowing sellers' prices in period 2 to adjust to the platform's choice of λ (given they are set after the level of λ is determined). It would be straightforward to extend our analysis by modeling the specific mechanism through which the platform would steer buyers towards risky or safe sellers, but that lies outside the scope of the current paper, which focuses only on the direction in which platform would like to steer buyers and the main factors that drive it.

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 incentivize exploration. 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 preferred 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 and multiple identical sellers of the safe product. This scenario applies, for example, when a new seller is granted a patent on a new product that has yet to be market tested. It also applies when a new seller enters with a product or service that is differentiated from existing competing sellers. For brevity, we will refer to the seller of the new (risky) product as the risky seller and the competing sellers of the established (safe) product as 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 implies

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

Improving signal accuracy (i.e., increasing the value of F through setting a higher λ) always increases the risky seller's expected profits in period 2. This reflects that 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. Note that while a higher λ also means the risky seller wins the market more often, this has no effect on its period-2 profits because for the marginal signal realization, the risky seller makes no profit.

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,[3] 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.

Now compare this to the platform's objective, which is to maximize its revenue from seller commissions. Based on the equilibrium prices charged by each type of seller in each period, the platform's profit is

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

[3] 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.

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 level of λ maximizing (6) is exactly the same way as the level maximizing (5).

The common level of exploration that the risky seller and the platform both want is determined by the first-order 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 equilibrium level of exploration, which is positive, is also the level of exploration preferred by the platform.*

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 interests of the platform and the risky seller coincide when it comes to setting λ.

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 4.2, 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 $(1-\beta)\,\pi_r(\lambda)$, while the platform's expected profit becomes

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

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, which we call the safe seller, and many identical sellers of the risky product, which we call the risky sellers. This captures the case of a monopolist of an established (safe) product that is facing entry from a potentially disruptive competitively supplied new (risky) product.

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

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

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

$$\pi_s'(\lambda) = \pi_r'(\lambda) > 0.$$

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

$$(1-\alpha)(1-\lambda)(u_s - \overline{u}_r) + \delta \pi_s(\lambda). \tag{8}$$

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.

Similar to Section 6, the platform's expected profit (from its share of revenue from each type of seller) is

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

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

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

Once again, the level of λ maximizing (9) is exactly the same way as the level maximizing (8).

Taking the derivative of (8) with respect to λ, the equilibrium level of exploration induced by the safe seller is the solution to

$$\pi_s'(\lambda) = \frac{1-\alpha}{\delta}(u_s - \overline{u}_r). \tag{10}$$

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 equilibrium level of exploration, which is positive, is also the level of exploration preferred by the platform. This common level is the same as the equilibrium level of exploration 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 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"). This means 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 (**??**) 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*

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

The condition that determines the equilibrium level of exploration (11) 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 (11) in the text of 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

$$\alpha\left(\lambda p_r^* + (1-\lambda)p_s^* + \frac{\delta}{1-\alpha}\left(\pi_r(\lambda) + \pi_s(\lambda) + c\right)\right), \tag{12}$$

which it maximizes with respect to λ. Consequently, the first-order condition determining the platform's preferred λ is also identical to (11).

Proposition 4 (Single seller of each type of product) *With one seller of the safe product and one seller of the risky product, the equilibrium level of exploration, which is positive, is also the level of exploration preferred by the platform. This common level is strictly higher than the equilibrium level of exploration 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 (11) 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 (11), to the one prevailing in Sections 3.2 and 3.3, which was determined by (7) and (10). The right-hand sides of the two equations are the same, but the left-hand side in (11) is twice the one in (10). 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 four 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 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 4.3 we allow the platform to invest in marketing to attract more buyers onto the platform and we show how this changes our benchmark results. Finally, 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.

4.1 Platform cares about 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, or it could be a shorthand for the fact the platform cares about buyer surplus due to the need to attract buyers in the future (i.e., reputational concerns). The equilibrium level of exploration remains unchanged in each of the four market configurations, but the platform's preference over λ 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, buyer surplus in each period 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 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 now cares about 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 exploration. 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. The platform's objective function (12) 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 now prefers a *lower* level of exploration than the equilibrium level induced by the two sellers, which remains unchanged and determined by (11). The more the platform cares about buyer surplus (i.e., the larger α_B), the larger the gap between the (excessive) equilibrium level of exploration and the level preferred by the platform.

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 equilibrium level of exploration, which equals zero, is insufficient from the platform's perspective.
(ii) If only one type of seller has market power, the equilibrium level of exploration, which is positive, is equal to the one preferred by the platform.
(iii) With a single seller of each type, the equilibrium level of exploration of the risky product is excessive from the platform's perspective.

Whether buyer surplus is increasing or decreasing in λ (which determines the direction of the bias) 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 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.[4] 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.[5]

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. Thus, the equilibrium level of exploration is $\lambda^* = 0$ if $\overline{u}_r - \frac{c_r}{1-\alpha_r} < u_s - \frac{c_s}{1-\alpha_s}$ and $\lambda^* = 1$ otherwise. In particular, if $c_r > c_s$ and $\alpha_r > \alpha_s$, then safe sellers make all sales in both periods and there is no exploration in equilibrium.

The platform's profit is

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

which 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 no exploration (i.e., $\lambda = 0$) if $\frac{\alpha_r c_r}{1-\alpha_r} < \frac{\alpha_s c_s}{1-\alpha_s}$ and full exploration (i.e., $\lambda = 1$) if $\frac{\alpha_r c_r}{1-\alpha_r} > \frac{\alpha_s c_s}{1-\alpha_s}$. Comparing with the equilibrium outcome, there is excessive (respectively, insufficient) exploration when c_r is small (respectively, large) enough relative to c_s or α_r is small (respectively, large) enough relative to α_s. In the remaining parameter range, where c_r is close enough to c_s and α_r is close enough to α_s, the equilibrium level of exploration is equal to that desired by the platform.

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

[4] 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.

[5] Corresponding to the assumption on (3) in the benchmark setting, we assume (13) below is concave in λ.

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 B.

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

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

In period 2, the risky seller's profit is

$$\pi_r(\lambda) \equiv (1-\alpha_r) \int_{\overline{u}_r + \frac{u_s - \overline{u}_r + \frac{c_r}{1-\alpha_r} - \frac{c_s}{1-\alpha_s}}{F(\lambda)}}^{u_H} \left(E_\lambda [u_r|x] - u_s + \frac{c_s}{1-\alpha_s} - \frac{c_r}{1-\alpha_r} \right) dG(x). \tag{13}$$

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

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

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

Meanwhile, the platform's profit is

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

which can be re-written as

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

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 wants more exploration relative to the level that maximizes the risky seller's total expected profit $\Pi_r(\lambda)$) if $\frac{\alpha_r c_r}{1-\alpha_r} > \frac{\alpha_s c_s}{1-\alpha_s}$ and decreasing in λ (so the platform wants less exploration than the risky seller) if $\frac{\alpha_r c_r}{1-\alpha_r} < \frac{\alpha_s c_s}{1-\alpha_s}$.

To understand this condition, 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 B. The following proposition summarizes the results of this section.

Proposition 6 (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.*
(i) With competitive sellers of both types, there is excessive exploration from the platform's perspective if

$$\frac{\alpha_r c_r}{1-\alpha_r} < \min\left\{\frac{\alpha_s c_s}{1-\alpha_s}, \alpha_r\left(\frac{c_s}{1-\alpha_s} - (u_s - \overline{u}_r)\right)\right\},$$

insufficient exploration if

$$\frac{\alpha_r c_r}{1-\alpha_r} > \max\left\{\frac{\alpha_s c_s}{1-\alpha_s}, \alpha_r\left(\frac{c_s}{1-\alpha_s} - (u_s - \overline{u}_r)\right)\right\},$$

and the right amount of exploration otherwise.
(ii) If at least one type of seller has market power, the equilibrium level of exploration is excessive from the platform's perspective if $\frac{\alpha_r c_r}{1-\alpha_r} < \frac{\alpha_s c_s}{1-\alpha_s}$ and insufficient 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.

4.3 Market expansion

Suppose now the platform is able to 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 its preferred choice of λ, the platform does not have to consider this effect, because it already chooses N optimally. As a result, the platform's ability to invest in buyer market expansion for period 2 will drive a wedge between the equilibrium level of exploration and the one preferred by the platform. 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 C.

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

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

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

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

The function $N(\lambda)$ in (15) is determined by the platform's profit maximization decision regarding

[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.

N given λ. Using the same steps as before, the platform's expected profit can be written as

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

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

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

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 (6) and its choice of N given λ is defined by (17). Using the envelope theorem, the first-order condition determining the platform's preferred λ is then

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

Comparing (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 $N'(\lambda)>0$. 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 C, the following proposition summarizes the main results for this section.

Proposition 7 (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, there is excessive exploration from the platform's perspective.

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.

4.4 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's and risky seller's objective function were maximized by the same choice of λ, 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, there is excessive exploration from the platform's perspective.*

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 new sellers have lower marginal costs or pay a lower revenue share to the platform relative to established sellers.* 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).

- *The platform can make costly investments to increase the number of participating buyers in response to the level of exploration determined by sellers.* 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.

- *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.

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.

Nevertheless, we also found two 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.*

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. A particularly interesting such consideration (which would provide a natural avenue for future research) is that buyers may (partially) attribute a negative experience with one seller in a particular product/service category to the platform. Thus, they may leave the platform altogether, which creates a negative spillover onto sellers in other product/service categories. In this context, when sellers in a given category consider how much buyer exploration to induce, they ignore the spillover from a negative buyer experience onto other product/service categories. On the other hand, arguably there is less of a spillover from a positive buyer experience—buyers continue using the platform. Consequently, this is another factor that can potentially lead to excessive exploration.

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 Appendix

6.1 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

$$p_r - p_s = \overline{u}_r - u_s, \tag{19}$$

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.[7]

The risky seller's total profits are

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

The risky seller must not want to slightly decrease p_r and get any number $\lambda' \in [\lambda, 1]$ of period-1 buyers, which means

$$\lambda\left(\left(1-\alpha\right)p_r - c\right) + \delta\pi_r\left(\lambda\right) \geq \lambda'\left(\left(1-\alpha\right)p_r - c\right) + \delta\pi_r\left(\lambda'\right)$$

for all $\lambda' \in [\lambda, 1]$. Since $\pi_r(\lambda)$ is concave, this is equivalent to

$$\left(\left(1-\alpha\right)p_r - c\right) + \delta\pi_r'\left(\lambda\right) \leq 0. \tag{20}$$

The risky seller must also not want to restrict the number of buyers receiving its period-1 subsidy below λ (i.e., keep price constant and get any number $\lambda' \in [0, \lambda]$ of period-1 buyers). This requires

$$\lambda\left(\left(1-\alpha\right)p_r - c\right) + \delta\pi_r\left(\lambda\right) \geq \lambda'\left(\left(1-\alpha\right)p_r - c\right) + \delta\pi_r\left(\lambda'\right)$$

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

for all $\lambda' \in [0, \lambda]$. Since $\pi_r(\lambda')$ is concave, this is equivalent to

$$((1-\alpha)p_r - c) + \delta\pi_r'(\lambda) \geq 0. \tag{21}$$

Note that this also implies

$$\lambda((1-\alpha)p_r - c) + \delta\pi_r(\lambda) \geq \delta\pi_r(0) > 0,$$

where the second inequality follows from assumption 2. Thus, the risky seller's equilibrium profits must be positive.

Together, (20) and (21) imply

$$(1-\alpha)p_r - c + \delta\pi_r'(\lambda) = 0, \tag{22}$$

so the equilibrium λ must maximize the risky seller's total profits.

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

$$-((1-\alpha)p_s - c) + \delta\pi_s'(\lambda) = 0. \tag{23}$$

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

Combining conditions (19), (22) and (23), the equilibrium λ and prices (p_r, p_s) must satisfy

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

7 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.

"The Economics of New Goods" Edited by Timothy F. Bresnahan and Robert J. Gordon, University of Chicago Press, 1997.

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.

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

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

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.

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 $E_\lambda [u_r|x] - p_r^2$. Thus, the buyer with signal x will purchase from the risky seller if and only if $x \geq \overline{u}_r + \frac{u_s - \overline{u}_r + p_r^2 - \frac{c}{1-\alpha}}{F(\lambda)}$. Thus, the risky seller's period-2 expected profit is

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

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.

*MIT Sloan School of Management
†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

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

Again, assumption (2) in the main paper implies that this profit is 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

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

Both profits are increasing in λ.

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

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

so the platform's profit expression (12) 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 Proof of Proposition 6

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

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

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, the platform's profit is

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

which can be re-written as

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

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 wants more exploration relative to the level that maximizes 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 (13) in the main paper and (B.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

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

We can also combine the equations above to obtain

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

The platform's profit as a function of λ (at equilibrium prices) is

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

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

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

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

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

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

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

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

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 equilibrium level of exploration is insufficient if $\frac{\alpha_r c_r}{1-\alpha_r} > \frac{\alpha_s c_s}{1-\alpha_s}$ and excessive if $\frac{\alpha_r c_r}{1-\alpha_r} < \frac{\alpha_s c_s}{1-\alpha_s}$.

C Proof of Proposition 7

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

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

Thus,

$$\frac{\partial \pi_s(\lambda, N)}{\partial \lambda} = \frac{\partial \pi_r(\lambda, N)}{\partial \lambda} \tag{C.2}$$
$$\frac{\partial \pi_s(\lambda, N)}{\partial N} = \frac{\partial \pi_r(\lambda, N)}{\partial N} + (1-\alpha)(u_s - \overline{u}_r), \tag{C.3}$$

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

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

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

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

Now consider the platform's problem. The platform's expected can be written as

$$\alpha\left((1-\lambda)(u_s - \overline{u}_r) + \frac{c}{1-\alpha}\right) + \frac{\delta\alpha}{1-\alpha}(\pi_s(\lambda, N) + (1+N)c) - C(N). \tag{C.5}$$

Given λ, the platform chooses N to maximize (C.5), 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 (C.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 and (C.2), the first-order condition determining the platform's preferred λ is identical to (18) in the main paper. Given the platform prefers the same level of λ as in the case with one risky seller and competitive safe sellers, but the safe seller now prefers an even higher level of λ, the excessive exploration result still holds.

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 (C.1), and the function $N(\lambda)$ is determined by the platform.

The platform's objective function is

$$\alpha\left(\lambda p_r^* + (1-\lambda)p_s^* + \frac{\delta}{1-\alpha}(\pi_s(\lambda, N) + \pi_r(\lambda, N) + (1+N)c)\right) - C(N), \tag{C.6}$$

where (p_r^*, p_s^*) are the two sellers' equilibrium prices in period 1 (which we will determine below). The platform maximizes (C.6) with respect to N and λ. 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*

$$\left.\frac{d\pi_s(\lambda, N(\lambda))}{d\lambda}\right|_{\lambda=\lambda^*} + \left.\frac{d\pi_r(\lambda, N(\lambda))}{d\lambda}\right|_{\lambda=\lambda^*} = \frac{1-\alpha}{\delta}(u_s - \overline{u}_r). \tag{C.7}$$

We now 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. From (C.6), using the envelope theorem and $p_s^* - p_r^* = u_s - \overline{u}_r$, we obtain that the first-order condition determining the platform's preferred λ is

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

Comparing (C.8) with (C.7) 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 excessive exploration result also holds in this case.