UNIVERSITÀ DEGLI STUDI FIRENZE
DISEI
DIPARTIMENTO DI SCIENZE PER L'ECONOMIA E L'IMPRESA

Working Papers - Economics

Recommendation Service as a Coordination Device among Peers

Elias Carroni, Giuseppe Pignataro, Alessandro Tampieri

Working Paper N. 06/2020

DISEI, Università degli Studi di Firenze
Via delle Pandette 9, 50127 Firenze (Italia) www.disei.unifi.it

Recommendation Service as a Coordination Device among Peers*

Elias Carroni[†]
Giuseppe Pignataro[‡]
Alessandro Tampieri[§]

September 2020

Abstract

In the classic beauty contest story of Morris and Shin (2002), the coordination game is set exogenously in the payoff function of agents. Our paper studies the existence of endogenous coordination in a context with one seller and many buyers. Due to imperfect information, the seller has an incentive to take a manipulative action to bias the product's quality. We show that a direct relationship between the seller and buyers does not induce them to coordinate to understand the product's value better. Coordination may occur if and only if a platform's intermediation is in place, as its recommendation to buyers unravels the beneficial effects of the others' documentation efforts. The platform's service works as a coordination device among peers. We finally model the combination of private and public signals with a platform's recommendation. The platform can have a preference for a *peer-review* system or an *individual learning* strategy, depending on the seller's ability to manipulate.

JEL codes: D42, D82, D83, L13, M37.

Keywords: Coordination Games, Recommendation, Platform, Manipulation, Bayesian Learning, Peer-Review.

*We would like to thank Francesca Barigozzi, Emilio Calvano, Stefano Colombo, Rahul Deb, Keisuke Hattori, Alessandro Lizzeri, Domenico Menicucci, David Myatt, Ignacio Monzón, Alessandro Pavan, Paolo Pin, Michelle Sovinsky, Yossi Spiegel, Salvatore Piccolo, Emanuele Tarantino, Giovanni Ursino, Jiekai Zhang as well as seminar participants at the 2019 Workshop on Institution, Individual Behavior and Economic Outcomes in Alghero, the 2018 EARIE conference in Athens, the XXXIII Jornadas de Economia Industrial in Barcelona, the 2018 ASSET Conference in Florence for helpful comments and suggestions. All errors remain our own.

[†]Department of Economics, University of Bologna, Email: elias.carroni@unibo.it.

[‡]Department of Economics, University of Bologna, Email: giuseppe.pignataro@unibo.it

[§]Department of Economics, University of Florence, Email: tamp79@gmail.com

1 Introduction

This paper analyzes the conditions that motivate the emergence of a coordination game among peers. Different social and economic environments, in reality, feature forms of coordination. Examples appear in almost all fields of decision-making such as financial-market interactions, firms' production, investment or innovation, voting or political decisions, e.g., Angeletos and Pavan (2004, 2007), Hellwig and Veldkamp (2009), Myatt and Wallace (2012, 2014), Vives (2011, 2017).

Agents are typically uncertain about the 'fundamental' and may have incentives to consider others' opinions as in the standard beauty contest (Keynes 1936). In case of complete information, both the fundamental and the 'beauty contest' terms of a typical coordination game coincide in the agent's payoff function and the interest in coordinating disappears. Instead, with imperfect information, agents receive public and private signals about the unknown fundamental value, and in turn, revise their beliefs. The diffusion of available information can reduce strategic uncertainty inducing a preference for coordination at the expense of the personal motive. In a seminal paper, Morris and Shin (2002) show that to the extent that agents respond to fundamental and coordination purposes, they have to choose actions based on both the expected (individual) value and the one set by the market, i.e., the average of all the agents' actions.

In the coordination-game literature, the two interacting motives have been usually set exogenously in the payoff function. Our contribution is to consider whether a coordination game among agents can endogenously appear due to agents' strategic choices. In particular, we argue that agents have incentives to coordinate only if a mediated interaction enters into the typical principal/agent relation.

Intermediaries (e.g., marketplaces) have a key role in online markets. Recent anecdotal evidence suggests an uptick in the role of peer recommendations and reviews to reduce the uncertainty about potential purchases. This real-world case is precisely an example in which people coordinate each other. For instance, according to the *2019 B2B Buyers Survey Report*, buyers spend much more time (effort) in online searching than before, thus increasing reliance on peers with multiple sources of information. Almost 92% of buyers give most credence to peer reviews and user-generated feedback to get informed about a product.[1]

Our setup describes the interaction between one seller and many buyers, where the fundamental value is the quality of the product. The seller may adopt deceptive practices, directed to change the perceived value of the product. Typical forms of deception involve prices, e.g., shrouding additional fees; or advertising, e.g., demand-shifting information. Buyers are *naive* in the sense that they cannot disentangle which part of the value they observe comes from the seller's manipulation. They may, however, exert effort by reading up and writing reviews that reflect their opinions on product quality.

[1] See *inter alia* http://view.ceros.com/g3-communications/dgr-2019-b2b-buyers-survey-report/p/1

We show that *without intermediation*, no buyer has individual incentives to make an effort to acquire information. The reason is that each buyer is not able, individually, to reduce the impact that the seller's deception may have on his perception of the product.

By contrast, deception can be reduced only by an aggregated activity of the buyers. This activity is induced by a mediator (a platform), which provides buyers with a purchase recommendation before receiving private and public signals. Such a recommendation is the main ingredient to ensure coordination among them. Therefore, *with intermediation*, coordination occurs since the platform can correctly anticipate buyers' average documentation effort. This evaluation is credibly possible in any online marketplace by, for instance, observing the repeated clicking behavior of each buyer. Then the platform can recommend purchases to buyers.

Amazon or Netflix are good examples of the approaches that a typical online mediator might take. Such platforms adopt a *collaborative filtering* process to make recommendations to clients even before buyers to start looking at individual items.[2] These recommendations are not based on products that are frequently bought together, but in particular on how buyers with similar purchase histories and interests behave.[3] Therefore, through a purchase recommendation, the platform may suggest the optimal individual demand, which best fits the agent's interest. Buyers internalize the effect of others' efforts and make an effort to correct the bias of manipulation in their perception of the quality.

Formally, once each buyer receives the personal recommendation and the market signals, a learning process takes place. We can identify three motives that explain a buyer's payoff: the *fundamental*, the *coordination*, and the *informative externality* motive. The first two are standard in coordination games. The *fundamental* motive highlights that the utility is higher (lower utility loss) as long as a buyer corrects more precisely the extent of the misperception. This term captures the distance between the effort and the quality of the product. The *coordination* motive instead identifies the impact of "higher-order beliefs", that is, the effect of the average effort on the individual documentation. Finally, the *informative externality* motive is due to the aggregate flow of information (through effort) that all buyers together collect and capture the influence that the (average) aggregation of information has on the perception of the fundamental value. It thus identifies the fundamental reason for the higher-order beliefs in the market, measuring the distance between the average aggregate effort and quality.

Next, we propose an additional ingredient that completes the functioning of a learning process in online marketplaces: the implementation of a *peer-review* system. In this case, buyers

[2] The *collaborative filtering* is a method of making automatic predictions or filtering about the users' interests by collecting preference and taste information from many users (collaborating). See for instance https://towardsdatascience.com/intro-to-recommender-system-collaborative-filtering-64a238194a26.

[3] Amazon claims its recommendation strategy as: "We determine your interests by examining the items you have purchased, items you have told us you own items you have rated, and items you have told us you like. We then compare your activity on our site with that of other customers, and using this comparison; we are able to recommend other items that may interest you." Amazon motivates this peer process on its Recommendations FAQ page: https://www.amazon.in/gp/help/customer/display.html?nodeId=201436180.

can fully observe others' efforts and choose how much information to share. Alternatively, the platform can obfuscate the news circulating among buyers, so that each one will learn only based on her private signal and the aggregate information given by the platform with the recommendation. We label this scenario as *individual learning*. This benchmark stylizes a real-world scenario in which posting an official comment on a service, product, or review is restricted under certain conditions, as made by Facebook or Instagram.

Although we do not explicitly model the platform's objective, we provide some insights on the platform's choice between *individual learning* and a *peer-review* system. In particular, given that the platform has interests to have both sides of the market on board, it may be optimal to induce an average effort between the optimal effort for buyers (positive) and the optimal one for the seller (zero). Allowing for a *peer review* may lead to insufficient or excessive learning from the buyers' viewpoint, depending on the seller's manipulative ability and the precision of each signal.

The remainder of the paper is organized as follows. Section 2 briefly surveys the related literature. Section 3 introduces the model setting. Then, in Section 4 and 5 we study the intermediation service provided by the platform and the buyers' learning process, respectively. Section 6 compare both *individual learning* and *peer-review* systems. Concluding remarks follow in Section 7.

2 Related literature

The paper first recalls the literature on the social value of information popularized by Morris and Shin (2002) and afterward developed by several contributions.[4] The game has a similar structure to the one proposed in the literature concerned with information sharing, see, e.g., Vives (1993), Angeletos and Pavan (2004, 2007), where the Gaussian-quadratic model and the linear solutions are common hypotheses.[5]

Different from the previous contributions, in this paper coordination endogenously emerges based on players' strategies. First, we show that a coordination game is possible to arise through a recommendation service offered by a platform that intermediates between a seller and many buyers. Second, we find a third, novel reason for coordination that adds to the fundamental and coordination motives. More precisely, internalizing others' efforts based on the service of the platform creates an additional *informative externalities* term of the buyer, who takes into account how much the aggregate flow of information is distant from the fundamental value.

The effect of the strategic complementarity or substitutability is similar to Angeletos and Pavan (2007) and Angeletos and Pavan (2007). This investigation line was even developed by

[4] See among others Colombo et al. (2014).

[5] A similar information structure is possible in a Lucas-Phelps island setting in Myatt and Wallace (2014).

Hellwig and Veldkamp (2009) and Hellwig and Veldkamp (2009) who investigate the combination of private and public information in coordinating settings. They show that adopting strategic complementarities in actions incentivizes the importance of *higher-order* beliefs. Such an equilibrium strategy was recently proposed and enriched in many settings with asymmetric information, including herding in networks (Battigalli et al., 2018; Bolletta and Pin, 2019), financial markets (Allen et al., 2006), business cycle models (Angeletos and La'O, 2009), and oligopolistic competition (Myatt and Wallace, 2015, 2018).

The literature slightly distinguishes between public and private information by looking at two exogenous signals. Instead, we believe that multiple information sources may enrich the design and constitute a fertile ground to develop. Therefore, each buyer receives a private signal correlated with all the other buyers' signals in our model.

Our paper is conceptually related to Burguet and Vives (2000) and Myatt and Wallace (2012) which characterize an information acquisition process of agents. In particular, Myatt and Wallace (2012, 2015) show how a costly acquisition process is possible as a combination of personal attention and the precision of the signal. In particular, agents have access to different information sources. The importance of each source depends on both the accuracy, i.e., the sender's noise, and the individual attention devoted to the information source, i.e., the receiver's noise.

In our setting, the information acquisition arises endogenously by a linear strategy solution due to a recommendation service of the platform as in reality. We propose a platform-mediated interaction between a seller and many buyers showing that the direct nonmediated communication between sellers and buyers does not allow for a coordination game among peers. This aspect is related to Fradkin (2017), who points out how marketplaces base their existence on the development of reviews. He suggests that "for any marketplace, informative reviews are a public good because writing reviews takes effort and has the potential to trigger retaliation" and "a second problem concerns the best manner in which to use review information throughout the platform. Importantly, these two choices are related because the incentives of reviewers depend on how the marketplace uses those reviews".

Note that the paper's purpose is not to model the optimization process of a platform developing a buyers' review system. This analysis is proposed by Vellodi (2018), who studies the impact of review reports in a model with no manipulation where a platform designs a rating system of heterogeneous sellers. He shows that full information transparency paradoxically works as an entry barrier and can cause adverse selection. In similar lines, Acemoglu et al. (2019) show how the platform may allow buyers to learn based on the ratings in a dynamic setting. They investigate the learning speed, disentangling the effect of collecting data from peer review compared to one of the rating systems. In our static model, the platform does not provide any rating system but comes up with a recommendation service on the product's quantity to buy. To some extent, our approach is complementary to theirs, as we give a rationale to the development of a *peer-review* system by directly studying the coordination problem among buyers, whose aggregate effort reduces the extent of manipulation.

3 The model

Consider the following one-shot game with one seller and n buyers. Buyer i's willingness to pay is denoted as v_i. Following Häckner (2000), we assume utility to be quadratic in consumption, represented by the demand D_i:

$$U_i = \left(v_i - p - \frac{D_i}{2} \right) D_i. \tag{1}$$

The optimal individual demand resulting from eq. (1) is linear. If the product is offered at price p, the net utility is maximized by the demand D_i solving the first-order conditions, i.e.:

$$\frac{\partial U_i}{\partial D_i} = 0 \Leftrightarrow D_i = v_i - p.$$

The seller decides upon a manipulative strategy to sell a product of value θ, which is private information. The seller's strategy consists in action $a \geq 0$, aiming at increasing the demand for the product, and a price p, aiming at extracting the surplus created by the manipulative action a. The manipulative strategy, hereafter simply *manipulation*, entails a quadratic cost:

$$\mathcal{C}(a) = \lambda a^2, \tag{2}$$

An increase in λ figures out a reduction of the manipulative ability of a seller. We define a seller who has a small (large) λ as highly (slightly) manipulative. Normalizing the marginal cost of production to zero without loss of generality, seller's profit writes as follows:

$$\pi = p \times D(a, p) - \mathcal{C}(a). \tag{3}$$

where $D(a, p)$ is the total demand.[6]

In case of perfect information, buyers observe the product value θ, so that $v_i = \theta$. In turn, each buyer i's willingness to pay corresponds to the product's real value, and all buyers will consume the same quantity with no possibility for the firm to manipulate, i.e., $a = 0$. In this case, the optimal price would be equal to $\theta/2$, and thus a buyer would receive a utility:

$$U_{PI} = \frac{3\theta^2}{8}, \tag{4}$$

where subscript PI stands for perfect information. Summing up, the total surplus of buyers will be given by $\sum_{i=1}^{n} U_{PI} = \frac{3n\theta^2}{8}$ whereas the seller gets the monopoly profits $\pi_{PI} = \frac{n\theta^2}{4}$.

Instead, under imperfect information, the seller may find it convenient to take a manipulative action to change the buyers' perception of the product value and make each valuation v_i

[6]It is necessary that $\lambda > n/4$ in $\mathcal{C}(a)$ to ensure concave profits.

higher than the real value of the product θ.[7] In turn, each buyer may have an incentive to exert effort, e_i, in reading reviews and sharing information with peers to render v_i closer to θ, thus making a more conscious purchase decision.

Hence, with imperfect information, v_i is generically different from θ and possibly idiosyncratic. It identifies the *perceived* value of the product as a function of the seller's manipulative strategy, and on the information each buyer gets. If the buyer does not exert any effort and just takes what the seller claims for granted, the willingness to pay of all buyers is homogeneous and equal to $v = \theta + a$. Exerting effort helps the buyer to make a more informed purchasing ex-ante, and gives $v_i = \theta + a - e_i$. Effort requires to face a quadratic cost, so that to the utility in eq. (1), we must subtract $-e_i^2$. The net utility is:

$$V_i = U_i - e_i^2. \tag{5}$$

In this sense, effort is inefficient as in Johnen and Somogyi (2019). It is exerted to correct *manipulation* for the firm, thus reducing the *perceived* utility, and restore the correct willingness to pay as much as possible closer to θ value as $e_i \leq a$. Note that the level of manipulation a adds to increase sales for any given level of the price, so that the total demand will be $D(a, p) = \sum_{i=1}^{n}(v_i - p)$.

Buyer's Information Structure. Each buyer i cannot observe the product quality even in case of no firm manipulation. The buyers have a common prior of the product value $\theta \sim \mathcal{N}(\bar{\theta}, \sigma_\theta^2)$ and each receives a private signal about θ, defined as $s_i = \theta + \epsilon_i$ where $\epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$, and a public signal, $z = a^*(\theta) + \omega$ where $\omega \sim \mathcal{N}(0, \sigma_\omega^2)$ which is based on the seller's level of manipulation.[8] In the case of mediated interaction, buyers also observe a purchase recommendation provided by the platform. It is not allowed in case of no mediation.

Conditional on θ, covariances are $cov[\omega, \epsilon_i] = 0$ and $cov[\theta, \epsilon_i] = 0$ for each buyer i. We assume the existence of the correlation among private signals. Formally the correlation coefficient between signals s_i and s_l, $\forall i \neq l$, is $\rho = \sigma_\omega^2/(\sigma_\omega^2 + \sigma_\epsilon^2)$ with covariance $cov[\epsilon_i, \epsilon_l] = \rho\sigma_\epsilon^2$. Further, ϵ_i are identically distributed with common precision τ_ϵ.[9] In this case, the signals' correlation induces an aggregate level of information, i.e., $\tilde{s} \sim \mathcal{N}\left(\bar{\theta}, \tilde{\sigma}^2\right)$ with $\tilde{\sigma}^2 = \sigma_\theta^2 + \sigma_\rho^2$, $\sigma_\rho^2 = (1 + (n-1)\rho)\sigma_\epsilon^2$ and $cov[\tilde{s}, s_i] = var[\tilde{s}] = \tilde{\sigma}^2$, while the related precisions are $\tilde{\tau} = \tau_\theta + \tau_\rho$ and $\tau_\rho = (1 + (n-1)\rho)\tau_\epsilon$ and the public precision is τ_ω. The precision-weighted signal average is thus a sufficient statistic for the signal $s_1, ..., s_n$ so $\mathbb{E}[\theta/\tilde{s}] = \bar{\theta} + \frac{\tau_\rho}{\tilde{\tau}}\left(\frac{1}{n}\sum_{i=1}^{n} s_i - \bar{\theta}\right)$.

[7] Seller's manipulation influences the perception of the product value, without giving any benefits to buyers. As a consequence, it induces a higher demand in comparison with the case in which the buyer observes the real product value.

[8] Although the use of normal distributions is standard in the literature of Bayesian learning, it determines a positive probability that prices and quantities are negative. The problem resolves through a proper choice of variances of the distributions and the parameters.

[9] We can define $\mathbb{E}[\theta/s_i] = \xi s_i + (1-\xi)\bar{\theta}$ where $\xi = \frac{\tau_\epsilon}{\tau_\epsilon + \tau_\theta}$ and $\mathbb{E}[s_i/s_{lm}] = \mathbb{E}[\theta/s_{lj}] = \xi\rho s_{lm} + (1-\xi\rho)\bar{\theta}$.

The non-mediated interaction. We present here the benchmark of a non-mediated interaction to show how the direct relationship between the seller and buyers do not create enough incentives to induce effort (and coordination) among buyers. As general as possible, the seller sets her strategy (action and price). Then each buyer observes the price and decides how much effort to exert to get information about the product. Finally, depending on the average effort of the n buyers, each buyer chooses how many units to consume.

The timing of the problem at hand is the following:

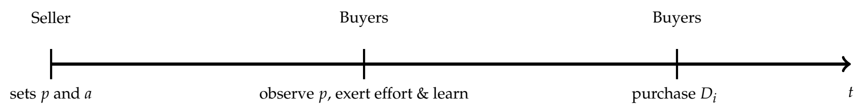

Figure 1: Direct relation between the seller and buyers

Such a simple formulation gives the following result:

Proposition 1. *When a seller directly interacts with the buyers, in equilibrium, buyers exert zero effort ($e_i = 0$ for all i) and the seller's strategies are:*

$$p^* = \frac{2\lambda\theta}{4\lambda - n}, \text{ and } a^* = \frac{n\theta}{4\lambda - n}.$$

Proof. See Appendix A.2. □

Proposition 1 is essential to highlight the fact that in our model, the intermediation of a platform is crucial to have incentives to exert effort. Absent the platform, the best for a single buyer would be to make no effort. Such a result is in line with the real-world observation that peer reviews and recommendation networks are often put in place by platforms that intermediate the interaction buyer-seller. Still, sellers' websites can hardly allow for reviews. Undoubtedly, this is because a non-mediated *peer-review* system is not credible.

In our model, the absence of buyers' effort is also intuitively linked to its nature and its impact on the seller's manipulative strategy. Since the individual effort only entails a cost without having the positive effect of reducing seller's manipulation, the buyer's optimal reply is to refrain from correcting it. Indeed, since the seller bases its strategy on aggregate demand, only the overall effort can induce it to reduce manipulation. However, to generate this total effort, the intervention of some intermediary is needed, and information can circulate among peers. We analyze both aspects in the next paragraph.

The mediated interaction. Now assume that a platform fully manages the interaction between the seller and the buyers. In particular, we consider the following setup. In the first stage, the platform communicates to the seller the demand function $D(a, p)$ and makes a purchase recommendation to each buyer, which takes the form of a quantity to buy D_i. Both the seller and the buyers make their decision based on the information the platform gives.

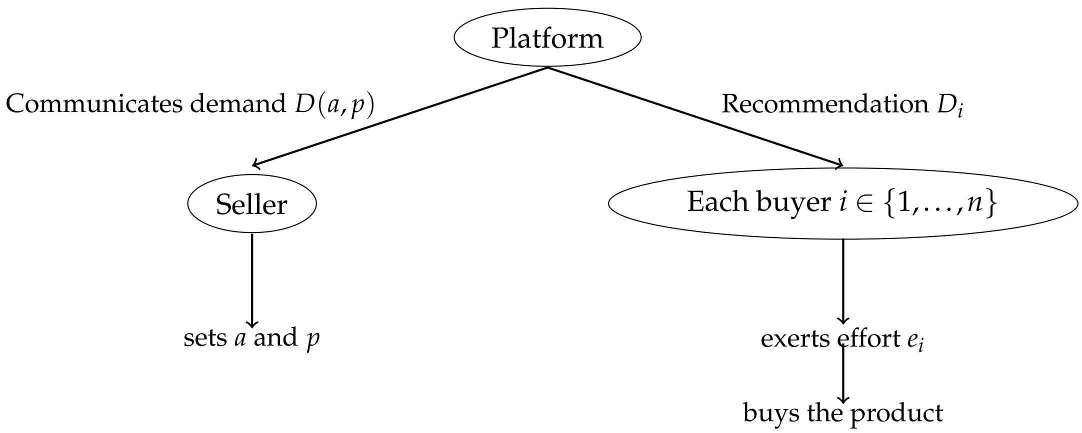

Figure 2: Intermediate relation between the seller and buyers

Importantly, both the demand function $D(a, p)$ and the recommendation to each buyer D_i will take into account the average aggregate effort $\bar{e} = \sum_{i=1}^{n} e_i/n$, which is correctly anticipated by the platform. For this reason, we can analyze the intermediation service and the learning process separately to understand exactly how a recommendation induces a coordination game among buyers. In particular, Section 4 is devoted to the analysis of the *intermediation service* with the average effort taken as given to understand what is the impact of a higher or lower average effort on the wellbeing of the seller, the buyers, and the platform. Then, Section 5 studies the learning process of consumers, which is different depending on whether the platform proposes a *peer-review* system or not.

4 Intermediation service

This section investigates the model when the effort exerted by buyers is taken as given. It helps in understanding how the buyers' effort influences the intermediation service of the platform. In the next section, we will enter into the details of the learning process.

Before presenting the results, let us express the nature of the intermediation. A purchasing recommendation from the platform is nothing more than algorithm-based advice linked to the buyers' average effort $\bar{e}$. These recommendations, in reality, like for Amazon or Facebook, are based on the similar purchase histories of buyers and the potential interests guiding their choices. Thus, the effort e is an all-catching representation of the clicking behavior to understand a product value. To be followed by each buyer, a recommendation should maximize the perceived utility before purchasing, after the learning process occurs. On the seller's side, the platform does not need to observe the quality of product θ to communicate to the seller the demand function $D(a,p) = n(I(\theta,\bar{e}) + a - p)$. We require that such demand function must have the following characteristics: *(i)* Purchase recommendations lead to the demand shifter $I(\theta,\bar{e})$ as a proxy of the product value θ and the buyers' (average) effort $\bar{e}$ and *(ii)* for each buyer, the recommendation is correct given the optimal strategy chosen by the seller. In the following proposition, we find how the recommendation is possible and that the market-clearing conditions realize based on the demand shifter.

Proposition 2. *When the platform mediates the interaction, the following holds:*

1. *The platform communicates to the seller the following demand function:*

$$D(a,p) = n(I(\theta,\bar{e}) + a - p)$$

   and recommends buyer i to buy a quantity $D_i = D(e_i,\bar{e}) = \frac{2\lambda(\theta-\bar{e})}{4\lambda-n} + \bar{e} - e_i$.

2. *The seller sets $p^* = \frac{I(\theta,\bar{e})+a^*}{2}$, where $a^* = \frac{nI(\theta,\bar{e})}{4\lambda-n}$.*

3. *The platform activity clears the market, so that $I(\theta,\bar{e}) = \theta - \bar{e}$.*

Proof. See Appendix A.3. □

The result in Proposition 2 must be read as follows. The platform communicates to the seller the demand function. For any given price, the demand decreases with the average effort, because the aggregate documentation activity reduces the manipulative ability of the seller to move the intercept of the demand function upwards. Intuitively, the more effort the buyer exerts to detect the product value; the less the seller can increase the intercept $I(\theta,\bar{e})$. Hence, for given λ, if buyers exert, on average, more effort in the documentation, the intercept of the demand function is reduced. The platform takes into account this when reporting the demand function to the seller so that there is less room to manipulate (a^* decreases), and the price p^* becomes lower.

On the buyer side, the recommendation is composed of two elements. The first one, $\frac{2\lambda(\theta-\bar{e})}{4\lambda-n}$, takes into account how increasing the average effort reduces the demand of the buyer who is exerting average effort. This captures the total impact that the collective documentation activity of all buyers has in reducing manipulation. The second one, $\bar{e} - e_i$, compares the average effort with the individual effort that each buyer will optimally decide to exert after having received the signals. From the individual viewpoint, the interpretation is straightforward. On average, peers exert an effort $\bar{e}$ to understand product value. Once each buyer receives a signal and optimally decides the documentation effort, her demand should be lower-than-average whenever she decides to exert more effort than the average ($e_i > \bar{e}$) to understand the product value. The opposite will occur whenever $e_i < \bar{e}$.

4.1 Buyer's utility

This section is devoted to describing the utility that the buyer realizes following the platform's recommendation. This utility differs from the *perceived* net utility V_i in eq. (5), since the true benefit obtained amounts to θ, and manipulation affects a buyer's perception without yielding any intrinsic benefit. Therefore, we compute the *realized* utility of individual i by considering the true value of the product θ:

$$U_{iR} = \left(\theta - p^* - \frac{D_i}{2}\right) D_i - e_i^2. \tag{6}$$

Hence, the average *realized* utility becomes:[10]

$$U_R = \frac{2\lambda^2\left(3\theta^2 - 9\bar{e}^2 - 2\bar{e}\theta\right) - n^2\bar{e}^2 + 2\lambda n\left(4\bar{e}^2 + \bar{e}\theta - \theta^2\right)}{(4\lambda - n)^2}. \tag{7}$$

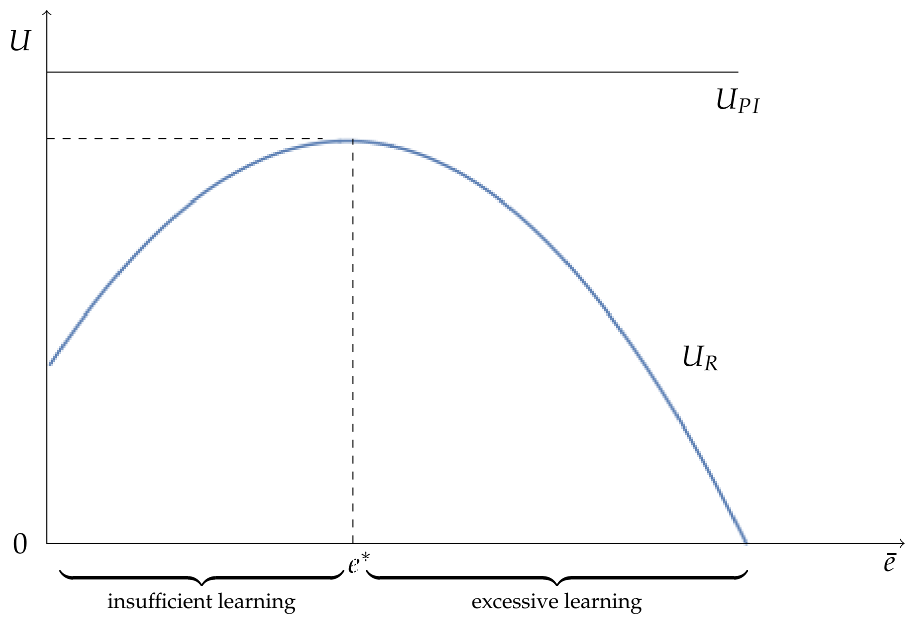

Figure 3: Consumer utility as a function of the average effort $\bar{e}$.

Some important considerations regard U_R. First, as observed in Figure 3, the utility U_R is concave in the average effort, which is useful to understand the value of the product better, but it also entails a quadratic cost. Moreover, the optimal average effort e^* depicted in Figure 3 is the one reported in the following Lemma.

[10]It identifies the *average realized* aggregate utilities of buyers.

Lemma 1. *U_R is maximized by setting an average effort equal to*

$$e^* = \max\left\{\frac{\theta\lambda(n-2\lambda)}{18\lambda^2+n^2-8\lambda n}, 0\right\}. \tag{8}$$

Proof. Notice that $\frac{\partial^2 U(\bar{e})}{\partial \bar{e}^2} = -\frac{2\left(18\lambda^2+n^2-8\lambda n\right)}{(4\lambda-n)^2} < 0$ for any $\lambda > n/4$. Therefore, the first-order conditions are sufficient for a maximum. Solving $\frac{\partial U(\bar{e})}{\partial \bar{e}} = 0$ with respect to $\bar{e}$ gives us the interior solution e^*. □

If the buyer's average effort is on the right of e^*, we say that the information environment leads to excessive learning, as buyers exert too much effort. Differently, if the average effort is on the left of e^*, the information environment leads to insufficient learning, as buyers would have something to gain to acquire more information about the product.

It is worth noticing the effect of a decrease in the seller's ability to shift demand, i.e., an increase in λ, on the buyer's optimal effort and maximal utility. First, it can be noticed that $\frac{\partial e^*}{\partial \lambda} < 0$ for any $e^* > 0$: a decrease in the manipulative ability of the seller maps into a lower effort required to maximize U_R. The demand shifts by a lower extent as λ increases, so that the need for documentation shrinks. As a consequence, the maximal utility $U_R(e^*)$ increases with λ.

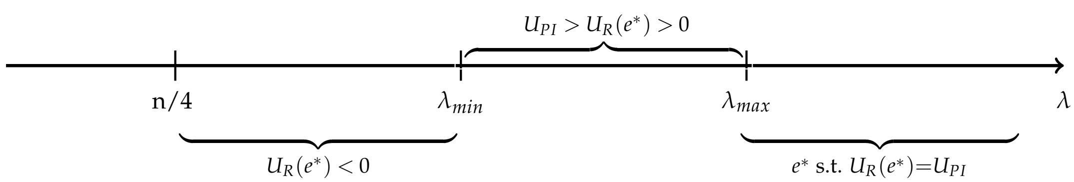

Figure 4: seller's type and the utility of recommendation.

(i) seller *highly manipulative* ($\lambda < \lambda_{min}$),

(ii) seller *moderately manipulative* ($\lambda_{min} < \lambda < \lambda_{max}$),

(iii) seller *poorly manipulative* ($\lambda > \lambda_{max}$).

Figure 4 shows the impact of the seller's type on the utility of the recommendation. For simplicity of exposition, a seller can be of three types. If $\lambda > \lambda_{max} \equiv \frac{n(\sqrt{22}-4)}{2}$, a seller is *poorly manipulative*. In these cases, the full-information utility can be reached, so that is sufficient to set e^* in such a way that $U_R(e^*) = U_{PI}$. Oppositely, if seller is *highly manipulative* (i.e., $\lambda < \lambda_{min} \equiv 2n/7$), the recommendation can never give a positive utility even when the effort is set to its optimal level, i.e., $U_R(e^*) < 0$.

4.2 Seller's profits

This section is devoted to the description of the impact of $\bar{e}$ on seller's profits. Plugging the optimal values $a^* = \frac{n(\theta-\bar{e})}{4\lambda-n}$ and $p^* = \frac{2\lambda(\theta-\bar{e})}{4\lambda-n}$ into the profit function in eq. (3), we obtain:

$$\pi^* = \frac{\lambda n(\theta - \bar{e})^2}{4\lambda - n}. \tag{9}$$

One can easily notice that $\frac{\partial \pi^*}{\partial \lambda} = -\frac{n^2(\theta-\bar{e})^2}{(n-4\lambda)^2} < 0$. This result is nothing surprising: the more costly it is for the seller to shift the demand, the lower the maximal profit that the seller will obtain. Therefore, for given product value, a highly manipulative seller makes higher profits than a poorly-manipulative seller. Moreover, π^* is decreasing in the average documentation effort, which reduces the possibility to shift demand. Therefore, if the objective is to maximize the seller's profits, the best would be to set $e_s^* = 0$.

5 Information acquisition

Here, we show how the platform recommendations generate an endogenous coordination game.

The learning process: a coordination game. Starting from eq. (1), a process of higher-order beliefs takes place as buyers observe the platform's recommendation $D_i = D(e_i, \bar{e})$ and the seller's price p. In particular, the platform's intermediation leads to the following result:

Proposition 3. *The game played among buyers is strategically equivalent to a quadratic-payoff coordination game, in which each buyer i seeks to minimize the following utility loss:*

$$\mathbb{E}[\mathcal{UL}_i(e_i, \bar{e}, \theta)] = -\kappa_1(e_i - \theta)^2 - \kappa_2(e_i - \bar{e})^2 - (1 - \kappa_1 - \kappa_2)(\bar{e} - \theta)^2$$

where $\kappa_1 = \frac{28\lambda^2+n^2-12\lambda n}{2(4\lambda-n)^2}$ *and* $\kappa_2 = \frac{14\lambda^2+n^2-8\lambda n}{(4\lambda-n)^2}$.[11]

Proof. See Appendix A.4. □

Proposition 3 states that the buyers' problem is similar to the minimization of a *utility-loss* function which is reminiscent of a *beauty contest* Bayesian coordination game *à la* Morris and Shin (2002). We can identify three motives that contribute to the quadratic payoff. The first term is the well-known *fundamental* motive and represents the action made by each buyer to understand the value of the product. In other words, the utility loss is lower as long as a buyer corrects more precisely the extent of misperception. The second term represents the *coordination* motive, well known in the literature of higher-order beliefs. Intuitively, each buyer could learn

[11] Note that the sufficient condition to get both κ_1 and κ_2 positive is when $\lambda > \frac{n}{14}(\sqrt{2} + 4)$.

due to social interactions with other people. The last term instead is an *informative externalities* motive due to the presence of higher-order beliefs among buyers. Such an information-sharing effect is a fundamental motive at the aggregate level.

We now study the buyers' learning process where buyers update their beliefs based on the private and public signals. We let the platform choose between two different learning environments. The first alternative is to allow for an *individual learning* process by avoiding the full sharing of information. In this case, each buyer can learn based only on her private signal and the public one. The second alternative is to offer a *peer-review* system, in which buyers receive the signals captured by the other peers and, consequently, can fully observe others' effort.

Formally, under *individual learning*, each buyer i is not able to understand the correlation coefficient between signals, i.e., $\rho = 0$, and chooses her optimal level of effort e_i conditional on the set $\{s_i, z\}$. Alternatively, with a *peer-review* system, each buyer i fully understands the correlation coefficient between signals, i.e., $\rho > 0$ and chooses her optimal level of effort e_i conditional on the set $\{s_i, \tilde{s}, z\}$.

5.1 Individual learning

Let us start with the case in which the platform allows buyers to evaluate her private signal plus the public signal when deciding their effort. The result reports in the following proposition.

Proposition 4. *Consider a case in which buyers engage in* individual learning *and let* $\Lambda = \frac{n}{4\lambda - n}$. *Buyer i who receives private signal s_i and public signal z optimally sets the following effort:*

$$e_i^{IL} = \alpha + \beta s_i + \delta z$$

where

$$\alpha = \bar{\theta}\beta(1-\mu)(1-\vartheta)(1+\kappa_2(1-\vartheta)), \qquad \beta = \frac{\kappa_1}{(1-\kappa_2\vartheta)}, \qquad \delta = \frac{\beta\mu(1-\vartheta)(1+\kappa_2(1-\vartheta))}{(1-\kappa_2)}$$

with $\vartheta = \tau_\epsilon/(\tau_\epsilon + \tau_\theta)$ and $\mu = \tau_\omega/(\tau_\theta + \Lambda\tau_w)$ represent the relative precision with respect to the prior of the public and the private signal, respectively.

Proof. See Appendix A.5. □

The weights α, β and δ say everything on the optimal level of effort of buyer i. Firstly, α weighs the effort component independent from the signals' realization. We label it as prior-based as it captures the weight of the prior on the optimal effort. Secondly, β gives the private signal's weight, which we label now on as the private-information part of the optimal effort. Finally, δ is the weight of the public signal. We will refer to this last as the public-information part of the optimal effort. The impact that signal precisions have on the optimal weights put by buyers to each signal in their learning process is stated in the following corollary.

Corollary 1. *Consider the optimal effort in Proposition 4. We have:*

(i) The prior-based part α increases in the precision of the prior τ_θ, while it decreases in the precision of the public τ_ω and the private signal τ_ϵ;

(ii) The private-information part β decreases in the precision of the prior τ_θ and increases in the precision of the private signal τ_ϵ, whereas it does not depend on the precision of the public signal τ_ω;

(iii) The public-information part δ increases in the precision of the public signal τ_ω, while it decreases in the precision of the private signal τ_ϵ. When the prior becomes more precise, i.e., τ_θ increases, buyers give more (less) weight to public information if the public signal is (not) sufficiently precise.

Proof. See Appendix A.6. □

Corollary 1 has a straightforward interpretation. The more precise a signal (or the prior), the more buyers tend to trust this source of information. This has two effects. On the one hand, a buyer will put more weight to a source of information in her documentation activity whenever it is more precise. This is captured by the fact that $\frac{\partial \alpha}{\partial \tau_\theta}$, $\frac{\partial \beta}{\partial \tau_\epsilon}$, and $\frac{\partial \delta}{\partial \tau_\omega}$ are all positive.

On the other hand, increasing the precision of a signal reduces the weight given to other sources of information, giving rise to the substitution effect. It is true, in particular, for the prior-based component, which decreases as public and private signals become more precise and weakly for the private-information part that falls in the precision of the prior. However, it is not affected by the precision of the public one. The only case in which the substitution effect reverses is between the precision of the prior and δ, as the weight given to the public signal increases with the precision of the prior when the public signal is precise enough.

Another important corollary of Proposition 4 is the following:

Corollary 2. *Consider the cases in which the signals or the prior are not precise. We have:*

$$\lim_{\tau_\epsilon \to 0} \alpha = \frac{\theta\kappa_1(\kappa_2+1)\left(\left(\Lambda^2-1\right)\tau_\omega+\tau_\theta\right)}{\Lambda^2\tau_\omega+\tau_\theta}, \lim_{\tau_\epsilon \to 0} \beta = \kappa_1 \textit{ and } \lim_{\tau_\epsilon \to 0} \delta = \frac{\theta\kappa_1(\kappa_2+1)\tau_\omega}{(1-\kappa_2)(\Lambda^2\tau_\omega+\tau_\theta)};$$

$$\lim_{\tau_\theta \to 0} \alpha = \lim_{\tau_\theta \to 0} \delta = 0 \textit{ and } \lim_{\tau_\theta \to 0} \beta = \frac{\kappa_1}{\kappa_2+1};$$

$$\lim_{\tau_\omega \to 0} \alpha = \infty, \lim_{\tau_\omega \to 0} \beta = \frac{\kappa_1(\tau_\theta+\tau_\epsilon)}{\tau_\theta+(1-\kappa_2)\tau_\epsilon} \textit{ and } \lim_{\tau_\omega \to 0} \delta = \frac{\kappa_1\tau_\theta\tau_\rho}{(1-\kappa_2)\left(\tau_\theta^2-\tau_\rho^2\right)}.$$

Corollary 2 is relevant as it highlights the inefficiencies of the learning process. In particular, given that the buyers evaluate their sources of information based on the signals received, they may end up taking into account information that has no precision at all. Corollary 2 shows that this is not the case for the public signal and the prior, as when they have zero precision, a buyer will ignore these sources of information. Differently, the buyer will always give some weight to the private signal, even if it is not providing any information (zero precision). Moreover, whenever a given signal's accuracy is zero, the weight given to the other sources of information is generically positive. An exception is when the prior precision is zero, which implies that the buyer gives zero weight to the prior and the public signal. This outcome is due to the substitution effect among the different sources of information, which reverses the relationship between the prior's precision and the public information part of the optimal effort.

5.2 Peer-Review system

Let us now consider the case in which the platform fully allows for information sharing among buyers. Each individual finds her private signal s_i, but also the aggregate level of information $\tilde{s}$ of the others based on the public signal when deciding her learning effort. The result reports in the following proposition.

Proposition 5. *Consider a case in which the platform offers a peer-review system. Buyer i who receives the private signal s_i, the aggregate level of information $\tilde{s}$ and the public signal z optimally sets the following effort:*

$$e_i^{PR} = \tilde{\alpha} + \tilde{\beta} s_i + \tilde{\delta} z,$$

where

$$\tilde{\alpha} = \frac{\theta\kappa_1}{1-\kappa_2} - \frac{\theta\tilde{\beta}(\Lambda\tau_\omega + \tau_\rho)}{\tau_\rho},\ \tilde{\beta} = \frac{\kappa_1^2\tau_\rho^2\left(\tau_\theta + \Lambda^2\tau_\omega\right)^2}{(1-\kappa_2)\Lambda^2\tau_\theta\tau_\omega\left(\left(\tau_\theta + \Lambda^2\tau_\omega\right)^2 - \tau_\rho^2\right)},\ \text{and}\ \tilde{\delta} = \frac{\tilde{\beta}\Lambda\tau_\omega}{\tau_\rho}.$$

Proof. See Appendix A.7. □

Similar to Corollary 1, the following result states the impact that signal precisions have on the optimal weights over each signal. It also says the effect of the correlation coefficient ρ, which plays a role only under a *peer-review* process.

Corollary 3. *Consider the optimal effort in Proposition 5. We have:*

(i) The prior-based part $\tilde{\alpha}$ increases in the precision of the prior τ_θ and of the public signal τ_ω, while it decreases in the precision of the private signal τ_ϵ;

(ii) *The private-information part $\tilde{\beta}$ decreases in the precision of the prior τ_θ and of the public signal τ_ω, while it increases in the precision of the private signal τ_ϵ;*

(iii) *The public-information part $\tilde{\delta}$ decreases in the precision of the prior τ_θ, whereas it increases in the precision of the private signal τ_ϵ and the public signal τ_ω.*

(iv) *The correlation coefficient ρ affects positively the weight given to the public and the private signals, reducing the weight given to the prior.*

Proof. See Appendix A.8. □

Corollary 3 is the information-sharing counterpart of Corollary 1. Intuitions are similar, so it is worth commenting only on the differences between the two, which are only on the effects of increasing the precision of the public signal on the prior-based and private-information component of the optimal effort. In particular, as long as the public signal gets more precise, a buyer will optimally put more (rather than less, as in Corollary 1) weight on the prior component and less weight on the private signal. The two results are related. When buyers do not share information, they learn individually, and the weight they put to the private signal is independent of the precision of the public signal. Consequently, when the precision of the public signal increases, only the weight given to the prior is affected negatively. Differently, under *peer-review*, buyers compare the information they receive privately with the information collected by the others and the public signal. Through these comparisons, the substitution effect arises unambiguously, so that a more precise public signal results in less weight given to the private one, *ceteris paribus*. All in all, this eventually results in a positive effect on the prior-based weight of $\tilde{\alpha}$.

In contrast with the case of *individual learning*, the correlation coefficient ρ plays a role here. Indeed, only when buyers share information with peers, they take into account the correlation among private signals. Point (iv) of Corollary 3 shows that the more correlated are private signals, the more buyers tend to give weight to the information they receive (public or private), thereby reducing the importance they give to their prior knowledge. This is not surprising. Considering the correlation among signals makes aggregated information more precise, making buyers more prone to rely on signals rather than prior experience.

To conclude this section, we show the weights when the precision of each signal tends to zero.

Corollary 4. *Consider the cases in which the precision of the signals is null. We have:*

$$\lim_{\tau_\epsilon \to 0} \tilde{\alpha} = \frac{\theta \kappa_1}{1 - \kappa_2} \text{ and } \lim_{\tau_\epsilon \to 0} \tilde{\beta} = \lim_{\tau_\epsilon \to 0} \tilde{\delta} = 0;$$

$$\lim_{\tau_\theta \to 0} \tilde{\alpha} = \lim_{\tau_\theta \to 0} \tilde{\beta} = \lim_{\tau_\theta \to 0} \tilde{\delta} = \infty;$$

$$\lim_{\tau_\omega \to 0} \tilde{\alpha} = \lim_{\tau_\omega \to 0} \tilde{\beta} = \infty \text{ and } \lim_{\tau_\omega \to 0} \tilde{\delta} = \frac{\kappa_1 \tau_\theta \tau_\epsilon}{(1-\kappa_2)\left(\tau_\theta^2 - \tau_\epsilon^2\right)}.$$

Again, results are similar to the ones in Corollary 2. The main difference is that if the private signal is maximally imprecise, buyers give zero weight to the public and private signals. Different from the case in which they do not share information, buyers can learn to a more significant extent about signals. Consequently, they also better learn about the lack of precision compared to the case of individual learning, finally leading them not to consider any information added to the prior in their learning activity. Notice that this effect is present only when the private signal's precision is zero, but it does not apply for the public one. Indeed, when the public signal's precision goes to zero, its weight is lower than the one of the other two sources of information, but it is still positive.

6 Peer-review system vs. Individual learning

This section investigates the potential effect of both learning environments. Given that the platform has interests to have both sides of the market on board, there are incentives to induce an average effort laying between the optimal effort for buyers (positive) and the optimal one for the seller (zero). Note that we do not model platform's objectives, but we just consider the platform as offering an environment that induces the lower but positive level of average effort. We will follow this rule in the following results, that compare the average effort under *individual learning* and *peer review* while varying the seller's manipulative ability.

The first and most natural case to study is how the precisions τ_θ, τ_ϵ, and τ_ω are symmetric. Note that we do not impose any restrictions on the information that contributes to the learning process. Indeed, the buyer's learning process will depend on prior and current information. The only caveat that all sources of information have the same variance and thus are all equivalently precise. This case is also critical because it can capture the relationship between the seller's manipulative nature and the effort emerging under the *individual learning* and the *peer-review* case. The results are in the following proposition.

Proposition 6. *Let $\tau_\theta = \tau_\epsilon = \tau_\omega$ and $\tilde{\lambda} \equiv \frac{\left(\sqrt{2}n+3\right)}{14}$. Two cases can arise:*

1. *if the seller is very manipulative ($\lambda < \tilde{\lambda}$), then the platform offers a* peer-review *system.*
2. *if the seller is moderately manipulative ($\lambda > \tilde{\lambda}$), then the platform designs an environment in which each buyer can only engage in* individual learning.

Proof. See Appendix A.9. □

The result describes the case when all sources of information have the same precision. Figure 5 helps to grasp the intuition. Recall that going from left to right means that the seller is less manipulative, as the cost of shifting demand increases. When the seller is very manipulative, there is a strong need for buyers to defend themselves from the seller's manipulative practice. When buyers are allowed to exchange information, this leads to high levels of effort, which decreases as long as manipulation is less important (λ increases). As a consequence, at some point, there exists a cutoff $\tilde{\lambda}$. Above this threshold, the equilibrium average effort e^{PR} in peer-review becomes zero, as manipulation is so moderate that there is no need for buyers to exert effort. The opposite occurs in the case of *individual learning*, i.e., e^{IL}. This is because when the seller is very manipulative ($\lambda < \tilde{\lambda}$), it is tough for each buyer to acquire information. Hence, buyers give up trying to understand the value of the product, and the average effort cannot be positive. As the level of manipulation decreases, a space for *individual learning* opens, and buyers exert more effort, and e^{IL} increases with λ.

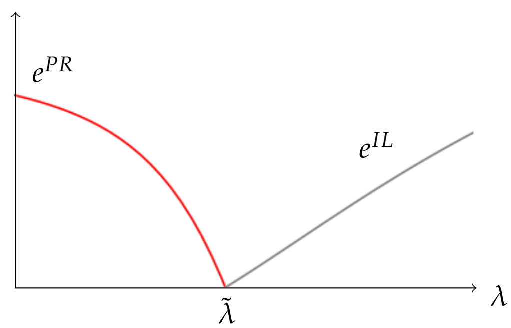

Figure 5: Effort levels when $\tau_\theta = \tau_\epsilon = \tau_\omega$. Effort resulting under *individual learning* (Grey curve) and the *peer-review* (Red Curve) as a function of λ with axes origin $(\lambda_{min}, 0)$.

In this context, the manipulation cost is crucial for the platform's choice. The platform can simply opt for a *peer review* process when the seller is highly manipulative, and for an *individual learning* when the seller is moderately manipulative.

Considering instead the extreme cases in which the prior is infinitely non-precise or infinitely precise, the results are slightly different. The first case is important to highlight what happens whenever agents only base their learning process on current information, given that the prior is not reliable. The first significant result is the following:

Lemma 2. *Assume that the prior is infinitely non-precise ($\tau_\theta \to 0$) and define $\hat{\lambda} \equiv \frac{(\sqrt{2}+3)n}{14}$. The platform will allow for a* peer-review *process if the seller is very manipulative ($\lambda < \hat{\lambda}$), whereas it designs an environment in which each buyer can only engage in individual learning.*

Proof. See Appendix A.10. □

The result here is similar to the one discussed in the previous section. The only difference is that a *peer-review* system always leads to a higher average effort as compared to *individual learning* (see Figure 6). The reason is that higher-order beliefs make buyers more prone to take

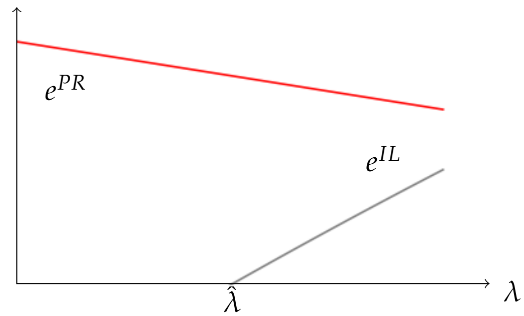

Figure 6: Effort levels when $\tau_\theta \to 0$. Top Panel: Effort resulting under *individual learning* (Grey curve) and a *peer-review* process (Red Curve) as a function of λ with axes Origin $(\lambda_{min}, 0)$.

into account information coming by other people. This effect is even more substantial as the importance of the current information prevails.

Moreover, the result that the average effort under a *peer-review* system decreases and under *individual learning* increases as the seller becomes less manipulative is confirmed. Differently, since ignoring the prior makes current information more critical than symmetric precisions, a *peer-review* system always leads to a higher effort.

Things drastically change when we consider the opposite extreme, i.e. when prior information becomes essential. Let us first state the results in the following Lemma.

Lemma 3. *Assume that the prior is infinitely precise ($\tau_\theta \to \infty$). In this case:*

1. *if the seller is very manipulative ($\lambda < \bar{\lambda}$), the platform will allow for a* peer-review *system;*
2. *if the seller is moderately manipulative ($\lambda > \bar{\lambda}$), the platform will allow for an* individual learning *system;*
3. *the choice of the platform leads to insufficient learning unless the seller is moderately manipulative.*

where $\bar{\lambda} \equiv \frac{\left(26+\sqrt{163-63\sqrt{5}}-2\sqrt{5}\right)n}{82}$.

Proof. See Appendix A.11. □

Figure 7 shows that the effort levels emerging in the two environments cross at zero when the seller is very manipulative. Then, they monotonically increase with a different slope. In particular, the effort is lower under *individual learning* when the seller is more manipulative, and the opposite is exact as the seller becomes less manipulative.

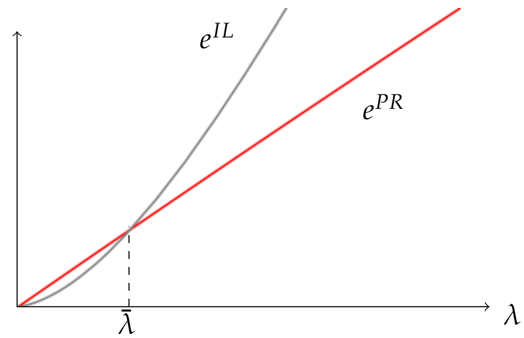

Figure 7: Effort levels when $\tau_\theta \to \infty$. Top Panel: Effort resulting under *individual learning* (Grey curve) and *peer-review* (Red Curve) as a function of λ with axes origin $(\hat{\lambda}, 0)$.

7 Concluding remarks

When a seller can inflate the perception of its product's value, and buyers face some documentation cost, manipulation can occur, and an individual buyer would never exert any effort. Indeed, to reduce the seller's incentives to manipulation, buyers need collective learning, which implies a coordination problem. Through their intermediation service, online marketplaces are often able to induce buyers to exert some documentation effort that would not be faced by an individual alone. Thus, this paper presents a tractable set up to give a rationale to the emergence of learning-by-peers online in response to possible manipulative strategies.

Our main result is that the recommendation service usually provided by online intermediaries is a device necessary to create a coordination game among buyers. As said, coordination is not possible in a direct relationship between the seller and the buyers as the latter cannot internalize others' efforts. In our story, the platform collects aggregate information about buyers' clicking behavior. Further, it can use it to (i) recommend an optimal purchase behavior to each consumer depending on the signals she receives, and (ii) inform the seller about the demand function. In this way, a coordination game among buyers emerges, as each individual may exert some positive effort to reduce *manipulation*.

The second take-home message is that the platform may not always have incentives to develop a *peer-review* system. Such a decision depends on the variation of the signal distributions and the seller's level of manipulation. Varying the precision of the signals received by each buyer may provide relevant insights into the dynamics of the information flow when the seller's manipulation increases. For example (among others), when all information sources have the same precision or when prior information plays no role, the platform would develop a *peer-review* system only if the seller is very manipulative.

References

Acemoglu, D., Makhdoumi, A., Malekian, A., Ozdaglar, A., 2019. Learning from reviews: the Selection effect and the Speed of Learning. Working Papers. URL: https://economics.mit.edu/files/17178.

Allen, F., Morris, S., Shin, H., 2006. Beauty contests and iterated expectations in asset markets. Review of Financial Studies 19, 719–752.

Angeletos, G., La'O, J., 2009. Noisy business cycles. NBER Macroeconomics Annual 24, 319–378.

Angeletos, G., Pavan, A., 2004. Transparency of information and coordination in economies with investment complementarities. American Economic Review P&P 94, 91–98.

Angeletos, G., Pavan, A., 2007. Efficient use of information and social value of information. Econometrica 75, 1103–1142.

Battigalli, P., Panebianco, F., Pin, P., 2018. Learning and selfconfirming equilibria in network games. arXiv preprint arXiv:1812.11775 .

Bolletta, U., Pin, P., 2019. Polarization when people choose their peers. Available at SSRN 3245800 .

Burguet, R., Vives, X., 2000. Social learning and costly information acquisition. Economic Theory 15, 185–205.

Colombo, L., Femminis, G., Pavan, A., 2014. Information acquisition and welfare. The Review of Economic Studies 81, 1438–1483.

Fradkin, A., 2017. Digital Marketplaces. Palgrave Macmillan UK. pp. 1–14.

Häckner, J., 2000. A note on price and quantity competition in differentiated oligopolies. Journal of Economic Theory 93, 233–239.

Hellwig, C., Veldkamp, L., 2009. Knowing what others know: Coordination motives in information acquisition,. Review of Economic Studies 76, 223–251.

Johnen, J., Somogyi, R., 2019. Deceptive Products on Platforms. Working Papers 19-13. NET Institute. URL: https://ideas.repec.org/p/net/wpaper/1913.html.

Keynes, J.M., 1936. The General Theory of Employment, Interest and Money. Independent Publishing Platform.

Morris, S., Shin, H.S., 2002. The social value of public information. American Economic Review 92, 1521–1534.

Myatt, D., Wallace, C., 2012. Endogenous information acquisition in coordination games. Review of Economic Studies 79, 340–374.

Myatt, D., Wallace, C., 2014. Central bank communication design in a lucas-phelps economy. Journal of Monetary Economics 63, 64–79.

Myatt, D., Wallace, C., 2015. Cournot competition and the social value of information. Journal of Economic Theory 158, 466–506.

Myatt, D., Wallace, C., 2018. Information use and acquisition in price-setting oligopolies. Economic Journal 128, 845–886.

Vellodi, N., 2018. Ratings design and barriers to entry. SSRN Working Papers URL: https://ideas.repec.org/p/net/wpaper/18-13.html.

Vives, X., 1993. How fast do rational agents learn? Review of Economic Studies 60, 329–347.

Vives, X., 2011. Strategic supply function competition with private information. Econometrica 79, 1919–1966.

Vives, X., 2017. Endogenous public information and welfare in market games. Review of Economic Studies 84, 935–963.

A Appendix

A.1 Concavity of profits - Section 3

The Hessian matrix is:

$$\mathcal{H} = \begin{pmatrix} -2n & n \\ n & -2\lambda \end{pmatrix},$$

whose first element is negative, and its determinant is

$$\det \mathcal{H} = n(4\lambda - n) > 0,$$

for $\lambda > n/4$.

A.2 Proof of Proposition 1.

Consider the optimal purchase choice of buyer i when a level of effort e_i has been exerted. Given e_i and the price offered by the seller p, buyer i sets D_i to maximize:

$$U_i = \left(v_i - p - \frac{D_i}{2} \right) D_i,$$

and it is maximized by an individual demand $D_i^* = v_i - p$. Going one step backward, the buyer sets her level of effort. Plugging D_i^* into the utility, and exploiting the fact that $v_i = \theta + a - e_i$, we get a maximal utility equal to:

$$U_i(D_i^*) = \frac{(\theta + a - e_i - p)^2}{2}, \tag{10}$$

which is decreasing in e_i. As a result, the optimal choice of a buyer is to exert zero effort, i.e., $e_i = 0$. Moving backward, anticipating the demand $D(a,p) = \sum_{i=1}^{n}(v_i + a - p) = (\theta + a - p)n$, the seller solves

$$\max_{a,p} \pi = p \times D(a,p) - \mathcal{C}(a) \tag{11}$$

The associated first order conditions of p and a are, respectively:

$$\frac{\partial \pi}{\partial p} = n(\theta + a - 2p) = 0, \tag{12}$$

$$\frac{\partial \pi}{\partial a} = np - 2\lambda a, \tag{13}$$

The second order conditions hold for $\lambda > \underline{\lambda}$, where $\underline{\lambda} \equiv \frac{n}{4}$ (see Appendix A.1 for further details). Rearranging eqs. (12) and (13), the optimal price and action are:

$$p^* = \frac{2\lambda\theta}{4\lambda - n}, \tag{14}$$

$$a^* = \frac{n\theta}{4\lambda - n}. \tag{15}$$

Given eqs. (14)-(15), demand and profits are:

$$D^* = \frac{2\lambda n\theta}{4\lambda - n}, \tag{16}$$

$$\pi^* = \frac{\lambda n\theta^2}{4\lambda - n}. \tag{17}$$

The zero-production corner solution is possible for each buyer when his demand is expected to be weak. This is assumed away by allowing for non-negative solutions as market-clearing prices cannot be negative. This completes the proof.

A.3 Proof of Proposition 2

Assume the platform to communicate to a seller a linear demand of the type $D(a, p) = n(I(\theta, \bar{e}) + a - p)$. As indicated above, this demand function has the following characteristics: *(i)* Purchase recommendations lead to the intercept $I(\theta, \bar{e})$ and *(ii)* for each buyer, the recommendation is correct given the optimal strategy chosen by each firm.

Let us proof the statement of Proposition 2 in six steps.

1. *Communication to the seller.* Assume the platform to communicate to the seller a demand function of the type $D(a, p) = n(I(\theta, \bar{e}) + a - p)$.

2. *Optimal seller's choice.* The seller sets p and a to maximize the profit in equation (3). Notice that p and a are strategic complements, since a is used by a seller to increase demand in response to any price and, consequently, the final aim is to increase the price. For given a the price that satisfy the FOC is:

$$\frac{\partial \pi}{\partial p} = n(I(\theta, \bar{e}) + a - 2p) = 0 \Leftrightarrow p(a) = \frac{I(\theta, \bar{e}) + a}{2}. \tag{18}$$

   In order to maximize profits it must hold

$$\frac{\partial \pi}{\partial a} = np - 2\lambda a \geq 0, \tag{19}$$

   with equality for interior solutions. It is easy to notice that $\frac{\partial \pi}{\partial a}|_{p(a)} = n\frac{I(\theta, \bar{e}) + a}{2} - 2\lambda a > 0$ whenever $nI(\theta, \bar{e}) + (n - 4\lambda)a > 0$, which is true for any a when $\lambda < n/4$. In those cases, manipulation is so inexpensive that it would be maximal $a^* = \bar{a}$ with $p(\bar{a})$. Differently, when $\lambda > n/4$, profits are concave and thus we can just use eq. (18) and considering (19) with equality to get:

$$a^* = \frac{nI(\theta,\bar{e})}{4\lambda - n} \quad \textit{and} \quad p^* = \frac{2\lambda I(\theta,\bar{e})}{4\lambda - n}. \tag{20}$$

3. *Optimal recommendation for any strategy of the seller.* Let us start by a buyer i who exerts effort to understand the value of the product. Assume that this buyer exerts effort equal to e_i. For given a and p, the correct recommendation that the platform can give to this buyer is to buy the quantity that maximizes her utility, which is

$$D_i = \theta + a - p - e_i. \tag{21}$$

4. *Optimal recommendation and aggregate demand of product.* Given the recommendation D_i that the platform gives to each buyer, the total quantity recommended for seller's product is the sum of all individual recommendations:

$$\begin{aligned} D &= \sum_{i=1}^{n}(\theta + a - p - e_i) \\ &= n(\theta + a - p) - \sum_{i=1}^{n} e_i \\ &= n(\theta + a - p - \bar{e}), \end{aligned}$$

where $\bar{e} = \sum_{i=1}^{n} e_i / n$ is the average effort exerted by buyers to discover product value.

5. *Market clearing.* In order for the market to clear, we need that D found in point 2 equalizes the demand $D(a,p) = n(I(\theta,\bar{e}) + a - p)$ communicated to the seller. Therefore, it must be that $I(\theta,\bar{e}) = \theta - \bar{e}$. Plugging this into $D(a,p) = n(I(\theta,\bar{e}) + a - p)$ and exploiting (20), we get:

$$a^* = \max\left\{\frac{n(\theta - \bar{e})}{4\lambda - n}, \bar{a}\right\}. \tag{22}$$

6. *Optimal individual recommendation.* Plugging (22) and (18) into the recommended quantity in eq. (21), we get the optimal recommendation depending on the average effort and the individual effort, i.e.:

$$D_i(e_i,\bar{e}) = \frac{2\lambda(\theta - \bar{e})}{4\lambda - n} + \bar{e} - e_i.$$

A.4 Proof of Proposition 3.

Assume $\lambda > n/4$, so that p^* and $D_i = D(e_i, \bar{e})$ are internal and substitute both into eq. (1) and taking into account that $I = \theta - \bar{e}$ yield the following:

$$U_i(e_i, \bar{e}, \theta) = \frac{\bar{e}^2 (4\lambda^2 + n^2 - 4\lambda n)}{2(4\lambda - n)^2} + e_i \left(\frac{\bar{e}(-16\lambda^2 - 2n^2 + 12\lambda n)}{2(4\lambda - n)^2} + \frac{\theta (4\lambda n - 16\lambda^2)}{2(4\lambda - n)^2} \right) + \frac{e_i^2 (-16\lambda^2 - n^2 + 8\lambda n)}{2(4\lambda - n)^2} + \frac{\bar{e}\theta (8\lambda^2 - 4\lambda n)}{2(4\lambda - n)^2} - \frac{2\theta^2 \lambda^2}{(4\lambda - n)^2}. \tag{23}$$

The utility above can be collected according to the individual and average learning effort and the quality, i.e., e_i, $\bar{e}$ and θ, as:

$$U_i(e_i, \bar{e}, \theta) = \vartheta_1 e_i^2 + \vartheta_2 \bar{e}^2 + \vartheta_3 \theta^2 + \vartheta_4 e_i \bar{e} + \vartheta_5 e_i \theta + \vartheta_6 \bar{e}\theta, \tag{24}$$

where the payoff function is quadratic in the action profile $e \equiv (e_i)_{i \in N} \in \mathbb{R}^N$ and quality $\theta \in \mathbb{R}$, and it is symmetric with respect to the permutations of other buyers. The parameters ϑ_1, ϑ_2, ϑ_3, ϑ_4, ϑ_5, ϑ_6 are constants and dependent of λ and n such that,

$$\vartheta_1 = -\frac{1}{2}, \tag{25}$$

$$\vartheta_2 = \frac{(2\lambda - n)^2}{2(4\lambda - n)^2}, \tag{26}$$

$$\vartheta_3 = \frac{2\lambda^2}{(4\lambda - n)^2}, \tag{27}$$

$$\vartheta_4 = -\frac{2\lambda - n}{4\lambda - n}, \tag{28}$$

$$\vartheta_5 = \frac{2\lambda}{4\lambda - n}, \tag{29}$$

$$\vartheta_6 = \frac{2\lambda(n - 2\lambda)}{(4\lambda - n)^2}. \tag{30}$$

In particular, ϑ_1, ϑ_4, and ϑ_5 are coefficients of e_i and contribute to derive the buyer i's best response. The coefficients ϑ_1 and ϑ_5 are negative, while ϑ_2 and ϑ_3 are positive. The parameters ϑ_4 and ϑ_6 have an ambiguous sign. This implies that, for a given market equilibrium, the buyers' learning game may exhibit strategic complementarity or substitutability both in $\vartheta_4 \gtrless 0$ and $\vartheta_6 \gtrless 0$. The variation is possible according to the evolution of the cost parameter λ and market size n. For instance, provided $\lambda > n/2$, it follows that $\vartheta_4 < 0$ and $\vartheta_6 > 0$. Constants ϑ_2 , ϑ_3 and ϑ_6 are coefficients of terms not including e_i, and they do not directly influence the buyer i 's best response. However, the effect of higher order beliefs plays a role in the interplay between individual effort e_i, the average effort $\bar{e}$ and the quality of the product θ. Such game is therefore strategically equivalently to a quadratic-payoff coordination game. The equivalence is possible by defining,

$$\kappa_1 = \frac{28\lambda^2 + n^2 - 12\lambda n}{2(4\lambda - n)^2}, \tag{31}$$

$$\kappa_2 = \frac{14\lambda^2 + n^2 - 8\lambda n}{(4\lambda - n)^2}, \tag{32}$$

and rearranging the equilibrium condition, each buyer i seeks to minimize,

$$\mathbb{E}[\mathcal{UL}_i(e_i, \bar{e}, \theta)] = -\kappa_1(e_i - \theta)^2 - \kappa_2(e_i - \bar{e})^2 - (1 - \kappa_1 - \kappa_2)(\bar{e} - \theta)^2, \tag{33}$$

where κ_1 and κ_2 are positive for $\lambda > \tilde{\lambda} \equiv \frac{n}{14}(\sqrt{2} + 4)$.

A.5 Proof of Proposition 4

The proof first shows the best response strategy $e_i^*(s_i, z)$ solving the first-order condition:

$$\mathbb{E}[e_i^*(s_i, z)] = \kappa_1 \mathbb{E}[\theta | s_i, z] + \kappa_2 \sum_{l=1}^{n} \mathbb{E}[e_l^*(s_l, z) | s_i, z], \tag{34}$$

where $z = a^*(\theta) + \omega$ and $a^*(\theta)$ is the optimal firm's action corresponding to $a^*(\theta) = \Lambda(\theta - \bar{e})$ with $\Lambda = \frac{n}{4\lambda - n}$ as in Subsection 4.2. Note that, in the case of individual learning, the platform does not allow for any *peer-review* process among users. Indeed, although the process is in higher-order beliefs as each buyer's learning effort strictly depends on others' effort, information cannot be transferred from one agent to another at all levels. From the properties of the Normal distribution, $y = \mathbb{E}[\theta | z] = \mu z + (1 - \mu)\bar{\theta}$ where $\mu = \tau_\omega / \tau_\theta + \Lambda^2 \tau_\omega$ and $\mathbb{E}[\theta | s_i, z] = \vartheta s_i + (1 - \vartheta) y$ as $\vartheta = \tau_\epsilon / (\tau_\epsilon + \tau_\theta)$ is linear in (s_i, z). By symmetry $e_i = \bar{e}$, we postulate a linear equilibrium strategy for buyer i of the form $e_i^*(s_i, z) = \alpha + \beta s_i + \delta z$ and for some $\alpha, \beta, \delta \in \mathbb{R}^+$ as:

$$\begin{aligned}
\mathbb{E}[e_i^*(s_i, z)] &= \kappa_1 \mathbb{E}[\theta | s_i, z] + \kappa_2 (\alpha + \beta \mathbb{E}[\theta | s_i, z] + \delta z) \\
&= (\kappa_1 + \kappa_2 \beta) \mathbb{E}[\theta | s_i, z] + \kappa_2(\alpha + \delta z) \\
&= (\kappa_1 + \kappa_2 \beta) \left(\vartheta s_i + (1 - \vartheta)\left(z\right) + (1 - \mu)\bar{\theta})\right) + \kappa_2(\alpha + \delta z),
\end{aligned}$$

Therefore, after some manipulations, we obtain that,

$$\alpha = \bar{\theta}\beta(1 - \mu)(1 - \vartheta)(1 + \kappa_2(1 - \vartheta)), \tag{35}$$

$$\beta = \frac{\kappa_1}{(1 - \kappa_2 \vartheta)}, \tag{36}$$

$$\delta = \frac{\beta\mu(1 - \vartheta)(1 + \kappa_2(1 - \vartheta))}{(1 - \kappa_2)}. \tag{37}$$

A.6 Proof of Corollary 1

Consider first the effects of the precisions on β. It is easy to notice that $\frac{\partial \beta}{\partial \tau_\omega} = 0$. Moreover, knowing that $\frac{\partial \vartheta}{\partial \tau_\epsilon} = \frac{\tau_\theta}{(\tau_\theta + \tau_\epsilon)^2} > 0$ and:

$$sign\left(\frac{\partial \beta}{\partial \tau_\epsilon}\right) = sign\left(\frac{\partial \vartheta}{\partial \tau_\epsilon}\right) = -sign\left(\frac{\partial \vartheta}{\partial \tau_\theta}\right).$$

Now consider α. It increases with the precision of the prior as:

$$\begin{aligned}
\frac{\partial \alpha}{\partial \tau_\theta} &= \frac{\alpha}{\beta}\frac{\partial \beta}{\partial \tau_\theta} - \frac{\alpha}{1-\mu}\frac{\partial \mu}{\partial \tau_\theta} - \frac{\alpha(2\kappa_2(1-\vartheta)+1)}{(1-\vartheta)(\kappa_2(1-\vartheta)+1)}\frac{\partial \vartheta}{\partial \tau_\theta} \\
&= \left(\underbrace{\frac{\kappa_1\kappa_2}{\beta(1-\kappa_2\vartheta)^2} - \frac{(2\kappa_2(1-\vartheta)+1)}{(1-\vartheta)(\kappa_2(1-\vartheta)+1)}}_{<0}\right)\alpha\underbrace{\frac{\partial \vartheta}{\partial \tau_\theta}}_{<0} - \frac{\alpha}{1-\mu}\underbrace{\frac{\partial \mu}{\partial \tau_\theta}}_{<0} > 0.
\end{aligned}$$

The precision of the public and the private signal reduce the effect of the prior, as:

$$\begin{aligned}
\frac{\partial \alpha}{\partial \tau_\epsilon} &= \frac{\alpha}{\beta}\frac{\partial \beta}{\partial \tau_\epsilon} - \frac{\alpha(2\kappa_2(1-\vartheta)+1)}{(1-\vartheta)(\kappa_2(1-\vartheta)+1)}\frac{\partial \vartheta}{\partial \tau_\epsilon} \\
&= \left(\underbrace{\frac{\kappa_1\kappa_2}{\beta(1-\kappa_2\vartheta)^2} - \frac{(2\kappa_2(1-\vartheta)+1)}{(1-\vartheta)(\kappa_2(1-\vartheta)+1)}}_{<0}\right)\alpha\underbrace{\frac{\partial \vartheta}{\partial \tau_\epsilon}}_{>0} < 0
\end{aligned}$$

and also, exploiting that $\frac{\partial \mu}{\partial \tau_\omega} = \frac{\tau_\theta}{(\tau_\theta + \Lambda\tau_\omega)^2} > 0$, we get:

$$sign\left(\frac{\partial \alpha}{\partial \tau_\omega}\right) = -sign\left(\frac{\partial \mu}{\partial \tau_\omega}\right) < 0.$$

Finally, let us consider the weight given to the public signal, δ. The last result can be used to show that $\frac{\partial \delta}{\partial \tau_\omega} > 0$ given that $sign\left(\frac{\partial \delta}{\partial \tau_\omega}\right) = sign\left(\frac{\partial \mu}{\partial \tau_\omega}\right) < 0$. Moreover, we have:

$$\begin{aligned}
\frac{\partial \delta}{\partial \tau_\epsilon} &= \frac{\delta}{\beta}\frac{\partial \beta}{\partial \tau_\epsilon} - \frac{\delta(2\kappa_2(1-\vartheta)+1)}{(1-\vartheta)(\kappa_2(1-\vartheta)+1)}\frac{\partial \vartheta}{\partial \tau_\epsilon} \\
&= \left(\underbrace{\frac{\kappa_1\kappa_2}{\beta(1-\kappa_2\vartheta)^2} - \frac{(2\kappa_2(1-\vartheta)+1)}{(1-\vartheta)(\kappa_2(1-\vartheta)+1)}}_{<0}\right)\delta\underbrace{\frac{\partial \vartheta}{\partial \tau_\epsilon}}_{>0} < 0,
\end{aligned}$$

and finally

$$\frac{\partial \delta}{\partial \tau_\theta} = \left(\underbrace{\frac{\kappa_1 \kappa_2}{\beta(1-\kappa_2\vartheta)^2} - \frac{(2\kappa_2(1-\vartheta)+1)}{(1-\vartheta)(\kappa_2(1-\vartheta)+1)}}_{<0} \right) \delta \underbrace{\frac{\partial \vartheta}{\partial \tau_\theta}}_{<0} + \frac{\delta}{\mu} \underbrace{\frac{\partial \mu}{\partial \tau_\theta}}_{<0}$$

has an ambiguous sign.

A.7 Proof of Proposition 5

The platform allows for a *peer-review* process among users. By taking into account a map from the signal space $(s_i, \tilde{s}, z)$ to the effort space $e_i(.)$, we can derive the best response $e_i(s_i, z)$ for all $\{s_1, ..., s_n, z\}$. The first-order condition is:

$$\mathbb{E}[\tilde{e}_i^*(s_i, z)] = \kappa_1 \mathbb{E}\left[\theta | s_1, ..., s_n, z\right] + \kappa_2 \sum_{l=1}^{n} \mathbb{E}\left[e_l(s_{lm}, z) | s_1, ..., s_n, z\right]. \tag{38}$$

Because $\tilde{y} = \mathbb{E}\left[\theta | s_i, \tilde{s},\ z\right]$ is linear in $(s_i, \tilde{s}, z)$, we can look at a linear solution. In particular, postulating a linear equilibrium strategy for buyer i of the form $\tilde{e}_i(s_i, z) = \tilde{\alpha} + \tilde{\beta} s_i + \tilde{\delta} z$. Equivalently, as we are searching for symmetric equilibrium such that any buyer $l \neq i$ has the same precision τ_ϵ and the same coefficients denoted by $(\tilde{\alpha}, \tilde{\beta}, \tilde{\delta})$ for the *candidate* effort function equal to $e_l(s_l, z)$, such that $\sum_{l=1}^{n} \mathbb{E}\left[e_l(s_{lm}, z) | \tilde{s}, z\right] = n\tilde{\alpha} + \tilde{\beta}\tilde{s} + n\tilde{\delta} z$. According to the information structure in Section 3, and similar to the proof of Proposition 4, we have again $a^*(\theta) = \Lambda(\theta - \bar{e})$ with $\Lambda = \frac{n}{4\lambda - n}$. Thus the public signal follows the information structure $z = \Lambda\left(\theta - \bar{e}\right) + \omega$. Substituting the expression for $\bar{e}$ according to the review process, and after some steps, we obtain $z = \tilde{\Lambda}\left(\theta - n\tilde{\alpha} - \tilde{\beta}\tilde{s}\right) + \omega$. Eq. (39) follows as:

$$\mathbb{E}[e_i(s_i, z)] = \kappa_1 \mathbb{E}[\theta / \tilde{s}, z] + \kappa_2 \left(n\alpha + \tilde{\beta}\tilde{s} + n\tilde{\delta} z\right). \tag{39}$$

From our Gaussian Information structure and applying the projection theorem:

$$\begin{pmatrix} \theta \\ \tilde{s} \\ z \end{pmatrix} \sim \mathcal{N}\left(\left(\begin{pmatrix} \bar{\theta} \\ \bar{\theta} \\ \bar{\theta} \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \sigma_\theta^2 & \sigma_\theta^2 \\ \sigma_\theta^2 & \sigma_\rho^2 & \tilde{\sigma}^2(\Lambda) \\ \sigma_\theta^2 & \tilde{\sigma}^2(\Lambda) & \tilde{\sigma}_\omega^2 \end{pmatrix} \right) \right),$$

where $\sigma_\rho^2 = (1 + (n-1)\rho)\sigma_\epsilon^2$ and $\tilde{\sigma}_\omega^2 = \Lambda^2 \sigma_\omega^2$, while $\tilde{\sigma}^2(\Lambda) = \sigma_\theta^2 + \tilde{\sigma}_\omega^2$. By projection theorem for normal random variables,

$$\mathbb{E}[\theta | \tilde{s}, z] = \bar{\theta} + \frac{\sigma_\theta^2 \tilde{\sigma}_\omega^2}{\Delta}(\tilde{s} - \bar{\theta}) + \frac{\sigma_\theta^2 \sigma_\rho^2}{\Delta}(z - \bar{\theta}), \tag{40}$$

where $\Delta = \left(\sigma_\rho^2 \tilde{\sigma}_\omega^2\right) - \left(\tilde{\sigma}^2(\Lambda)\right)^2$ is the variance-covariance correlation between z and $\tilde{s}$. Expressing (40) as a function of precisions τ_θ, τ_ω and τ_ϵ,[12]

$$\mathbb{E}[\theta|\tilde{s},z] = \bar{\theta} + \frac{\frac{1}{\tilde{\tau}_\omega}(\tilde{s}-\bar{\theta}) + \frac{1}{\tau_\rho}(z-\bar{\theta})}{\tau_\theta\left(\frac{(\tau_\theta+\Lambda^2\tau_\omega)^2-\tilde{\tau}_\epsilon^2}{\tau_\rho^2(\tau_\theta+\Lambda^2\tau_\omega)^2}\right)}.$$

Plugging the aggregate level of information $\tilde{s}$ and substituting $\mathbb{E}[\theta|\tilde{s},z]$ into (39) ,

$$\mathbb{E}[e_i(s_i,z)] = \kappa_1\left(\bar{\theta} + \frac{\frac{1}{\tilde{\tau}_\omega}(\tilde{s}-\bar{\theta}) + \frac{1}{\tilde{\tau}_\epsilon}(z-\bar{\theta})}{\tau_\theta\left(\frac{(\tau_\theta+\Lambda^2\tau_\omega)^2-\tau_\rho^2}{\tau_\rho^2(\tau_\theta+\Lambda^2\tau_\omega)^2}\right)}\right) + \kappa_2\left(n\alpha + \tilde{\beta}\tilde{s} + n\tilde{\delta}z\right).$$

Therefore after several manipulations, the solutions for $\tilde{\alpha}, \tilde{\beta}, \tilde{\delta} \in \mathbb{R}$ follow as,

$$\tilde{\alpha} = \bar{\theta}\kappa_1\left(\frac{(1+(n-1)\rho)\tau_\epsilon\left(\tau_\theta+\Lambda^2\tau_\omega\right)^2\left((1+(n-1)\rho)\tau_\epsilon+\Lambda^2\tau_\omega\right)}{\Lambda^2\tau_\theta\tau_\omega(1-\kappa_2)\left((1+(n-1)\rho)\tau_\epsilon)+\Lambda^2\tau_\omega+\tau_\theta\right)\left(\Lambda^2\tau_\omega+\tau_\theta(1+(n-1)\rho)\tau_\epsilon\right)} - 1\right) \tag{41}$$

while the weights of the private and public signals in the equilibrium solution, respectively, $\tilde{\beta}$ and $\tilde{\delta}$ are:

$$\tilde{\beta} = \frac{\kappa_1^2(1+(n-1)\rho)^2\tau_\epsilon^2\left(\tau_\theta+\Lambda^2\tau_\omega\right)^2}{(1-\kappa_2)\Lambda^2\tau_\theta\tau_\omega\left(\left(\tau_\theta+\Lambda^2\tau_\omega\right)^2-(1+(n-1)\rho)^2\tau_\epsilon^2\right)}, \tag{42}$$

$$\tilde{\delta} = \frac{\kappa_1(1+(n-1)\rho)\tau_\epsilon\left(\tau_\theta+\Lambda^2\tau_\omega\right)^2}{\tau_\theta(1-\kappa_2)\left((1+(n-1)\rho)\tau_\epsilon)^2-\left(\tau_\theta+\Lambda^2\tau_\omega\right)^2\right)}. \tag{43}$$

A.8 Proof of Corollary 3

Let us compute the derivatives of $\tilde{\alpha}, \tilde{\beta}$ and $\tilde{\delta}$ with respect of precisions. We have:

[12] Here let (X_1, X_2) simply be a random vector whose distribution is normal with $\mu_i = \mathbb{E}[X_i]$ and $C_{ij} = cov(X_i, X)$ $i,j = 1,2$. Then $\mathbb{E}[X_1|X_2] = \mu_2 + C_{21}C_{11}^{-1}(X_1-\mu_1)$.

$$\frac{\partial\tilde{\alpha}}{\partial\tau_\omega} \propto sign\left(\frac{1}{\tau_\theta - \kappa_2\tau_\theta}\right) > 0,$$
$$\frac{\partial\tilde{\alpha}}{\partial\tau_\theta} \propto sign\left(\frac{1}{(1-\kappa_2)}\right) > 0,$$
$$\frac{\partial\tilde{\alpha}}{\partial\tau_\epsilon} \propto sign\left(-\frac{\tau_\rho\Lambda}{(1-\kappa_2)}\right) < 0,$$
$$\frac{\partial\tilde{\beta}}{\partial\tau_\omega} \propto sign\left(-\frac{\tau_\epsilon(1+\Phi)^2}{(1-\kappa_2)}\right) < 0,$$
$$\frac{\partial\tilde{\beta}}{\partial\tau_\epsilon} \propto sign\left(\frac{1}{(1-\kappa_2)}\right) > 0,$$
$$\frac{\partial\tilde{\beta}}{\partial\tau_\theta} \propto sign\left(-\frac{\tau_\epsilon(1+\Phi)^2}{(1-\kappa_2)}\right) < 0,$$
$$\frac{\partial\tilde{\delta}}{\partial\tau_\epsilon} \propto sign\left(\frac{1}{(1-\kappa_2)}\right) > 0,$$
$$\frac{\partial\tilde{\delta}}{\partial\tau_\omega} \propto sign\left(\frac{1}{(1-\kappa_2)}\right) > 0,$$
$$\frac{\partial\tilde{\delta}}{\partial\tau_\theta} \propto sign\left(\frac{1}{(\kappa_2-1)}\right) < 0,$$

where $\Phi = (n-1)\rho$ and $\kappa_2 < 1$. Now, consider the effect of the correlation coefficient ρ. We get:

$$\frac{\partial\tilde{\alpha}}{\partial\rho} \propto sign\left(\frac{1}{(\kappa_2-1)}\right) < 0,$$
$$\frac{\partial\tilde{\beta}}{\partial\rho} \propto sign\left(\frac{1}{(1-\kappa_2)}\right) > 0,$$
$$\frac{\partial\tilde{\delta}}{\partial\rho} \propto sign\left(\frac{1}{(1-\kappa_2)}\right) > 0.$$

A.9 Proof of Proposition 6

Let $\tau_\theta = \tau_\epsilon = \tau_\omega = \tau$. Plugging into e^{PR} and e^{IL}, we get:

$$e^{PR} = \frac{4\theta\kappa_1\left(\left(\Lambda^2\tau+\tau\right)^2 + \frac{1}{2}(n+1)\left(\left(\Lambda-1\right)\left(\Lambda^2\tau+\tau\right)^2 - \frac{1}{2}(n+1)\tau^2\right)\right)}{\tau^2\left((n+1)\left(2\Lambda\left(\Lambda^2+1\right)^2\kappa_1 + (\kappa_2-1)(n+1)\right) - 4\left(\Lambda^2+1\right)^2(\kappa_2-1)\right)},$$

and

$$e^{IL} = \frac{4+8\Lambda-\theta\kappa_1(\kappa_2((\kappa_2+4)\Lambda+4))}{\kappa_2\Lambda(\kappa_1+2\kappa_2-6)+2(\kappa_1+2)\Lambda+2(\kappa_2-3)\kappa_2+4}.$$

The following results occur:

1. $\frac{\partial e^{PR}}{\partial \lambda} < 0$ when $\lambda > \lambda_{min}$. Notice also that $e^{PR} = 0$ when $\lambda = \tilde{\lambda} \equiv \frac{(\sqrt{2}n+3)}{14}$; while $e^{PR} > 0$ when $\lambda = \lambda_{min}$. As a consequence, if buyers are allowed to share information, this choice leads to a positive level of effort only if the seller is manipulative enough (λ is small).

2. $\frac{\partial e^{IL}}{\partial \lambda} > 0$ when $\lambda > \lambda_{min}$. Moreover, $e^{IL} = 0$ when $\lambda = \tilde{\lambda}$; while $e^{IL} < 0$ when $\lambda = \lambda_{min}$. As a consequence, if buyers are not allowed to share information, this choice leads to a positive level of effort only if the seller is not too manipulative (λ is high).

Hence, the platform can only opt for a *peer-review* when the seller is highly manipulative and an *individual learning* process when is moderately manipulative.[13]

A.10 Proof of Lemma 2

We compute the limits of the two efforts with τ_θ at zero:

$$\lim_{\tau_\theta \to 0} e^{PR} = \frac{2\theta (n - 2\lambda)}{n}$$

and

$$\lim_{\tau_\theta \to 0} e^{IL} = \frac{\theta \left(28\lambda^2 + n^2 - 12\lambda n\right)}{4\lambda^2}.$$

The following results occur:

1. $\frac{\partial \lim_{\tau \to 0} e^{PR}}{\partial \lambda} < 0$ when $\lambda > \lambda_{min}$. Then, $e^{PR} > 0$ when $\lambda = \lambda_{max}$. As a consequence, if the platform designs an environment in which buyers are allowed to share information, this choice will always lead to a positive level of effort.

2. $\frac{\partial \lim_{\tau \to 0} e^{IL}}{\partial \lambda} > 0$ when $\lambda > \lambda_{min}$. Then, $e^{IL} = 0$ when $\lambda = \hat{\lambda}$; while $e^{IL} < 0$ when $\lambda = \lambda_{min}$. As a consequence, if the platform designs an environment in which buyers are not allowed to share information, this choice leads to a positive level of effort only if the seller is not too manipulative (λ is high)

Thus, the platform can only opt for a *peer-review* process when the seller is highly manipulative and for *individual learning* when is moderately manipulative.

A.11 Proof of Lemma 3

We compute the limits of the two efforts with τ_θ at infinity:

$$\lim_{\tau_\theta \to \infty} e^{PR} = \frac{\theta \left(28\lambda^2 + n^2 - 12\lambda n\right)}{4\lambda^2}$$

[13] Notice that we obviously restrict our attention to situations in which the average effort is positive. Hence, whenever the average effort falls below zero, we do not consider this as a viable alternative for the platform when designing the environment for information acquisition.

and

$$\lim_{\tau_\theta \to \infty} e^{IL} = \frac{\theta \left(28\lambda^2 + n^2 - 12\lambda n\right) \left(46\lambda^2 + 3n^2 - 24\lambda n\right)}{2(4\lambda - n)^4}.$$

The following occurs:

1. $\lim_{\tau_\theta \to \infty} e^{PR} = \lim_{\tau \to 0} e^{IL}$ and $46\lambda^2 + 3n^2 - 24\lambda n > 28\lambda^2 + n^2 - 12\lambda n$ when $\lambda \in [\lambda_{min}, \lambda_{max}]$. Therefore, it is necessary and sufficient that $\lambda > \hat{\lambda}$ in order for the two efforts to be positive.

2. Moreover, when $\lambda = \hat{\lambda}$, $\lim_{\tau \to \infty} e^{PR} = \lim_{\tau \to \infty} e^{IL}$=0.

3. $\frac{\partial^2 \lim_{\tau \to \infty} e^{PR}}{\partial \lambda^2} < \frac{\partial^2 \lim_{\tau \to \infty} e^{IL}}{\partial \lambda^2}$. Therefore, there exists a cutoff value of λ above which $e^{IL} > e^{PR}$ and below which the opposite is true. The cutoff is the value of λ such that $\lim_{\tau_\theta \to \infty} e^{PR} = \lim_{\tau \to 0} e^{IL} > 0$ which is $\bar{\lambda} \equiv \frac{\left(26 + \sqrt{163 - 63\sqrt{5}} - 2\sqrt{5}\right) n}{82}$.