Seeding the Herd: Pricing and Welfare Effects of Social Learning Manipulation

Li Chen
Samuel Curtis Johnson Graduate School of Management · Cornell University
li.chen@cornell.edu

Yiangos Papanastasiou
Haas School of Business · University of California, Berkeley
yiangos@haas.berkeley.edu

This paper is motivated by the recent emergence of various interference tactics employed by sellers attempting to manipulate social learning. We revisit the classic model of observational social learning, and extend it to allow for (i) asymmetric information on product value between the seller and the consumers, and (ii) the ability of the seller to "seed" the observational learning process with a fake purchase, in an attempt to manipulate consumer beliefs. We examine the interaction between social learning manipulation and equilibrium market outcomes, as well as the impact of anti-manipulation measures aimed at detecting and punishing misconduct. The analysis yields three main insights. First, we show that increasing the intensity of anti-manipulation measures can have unintended consequences, often inducing higher levels of manipulation as well as higher equilibrium prices. Second, we find that although measures of high intensity can completely deter misconduct, such measures do not lead to any improvement in either seller or consumer payoffs, relative to the case where no measures are present. Third, we demonstrate that in many cases, measures of intermediate intensity can leverage seller manipulation to simultaneously improve both seller and consumer payoffs.

Key words: social learning, seller manipulation, applied game theory

1. Introduction

The recent paradigm shift from offline to online shopping has turned the consumer's purchase decision into a much more "social" process, one that is heavily linked to the purchases and experiences of other consumers. The typical modern-day consumer often does not recognize a need until this is brought to her attention by a blogger's article or a social-media post by one of her online "friends." Once she recognizes a need, she searches for the best product/solution by consulting best-seller lists and relying on product-display algorithms that use the searches and purchases of other consumers as an input. And when she narrows down her search, she evaluates the alternatives by reading her peers' comments and product reviews.

As the influence of various forms of *social learning* on consumers' purchase decisions has grown steadily over the last two decades, so has the incentive for sellers to engage in interference tactics in order to attract more consumers to their product. Indeed, reports in popular press now regularly document sellers' attempts to manipulate social learning in a variety of ways, including paying social-media "influencers" (i.e., consumers with many social connections) to publicly endorse their products (e.g., *Guardian* 2017, *BBC* 2017), faking purchase transactions to boost sales performance and product rankings (e.g., *Forbes* 2015, *The Wall Street Journal* 2015), and soliciting fake reviews to improve perceptions of consumer satisfaction (e.g., *Fortune* 2019, *Forbes* 2019; see also Mayzlin et al. 2014, Luca and Zervas 2016), among others.

In response to the growing awareness and concerns regarding sellers' manipulation tactics, defensive measures have been implemented at two levels. At the platform level, companies such as Alibaba, Amazon, and Instagram have undertaken various initiatives in order to detect and punish sellers' attempts to manipulate social learning (e.g., see *Financial Times* (2016) for Alibaba measures against fake purchases; *Fortune* (2016) for Amazon measures against fake reviews; *The Wall Street Journal* (2017) for Instagram measures against fake endorsements). At the government level, the US Federal Trade Commission has recently intensified its involvement by bringing cases against sellers allegedly engaged in manipulation tactics, and by issuing commercial-partnership disclosure guidelines and complaints procedures (see Federal Trade Commission (2019) for the first case against fake reviews; Federal Trade Commission (2016) for the first case against fake endorsements; Federal Trade Commission (2019) for guidelines and procedures).

Interestingly, despite the perceived importance of social learning manipulation, as well as the significant amounts of resources invested in combating seller interference, our understanding of this phenomenon and its basic implications remains limited. The goal of this paper is to enhance this understanding, by focusing on two first-order questions: First, what is the relationship between seller manipulation and market outcomes such as price, seller profit, and consumer surplus? Second, what is the impact of defensive measures on seller manipulation, and by extension on the aforementioned market outcomes?

As a first attempt to investigate these questions, we revisit the classic model of observational social learning first introduced by Banerjee (1992) and Bikhchandani et al. (1992), making two modifications: (i) the seller and the consumers are asymmetrically informed about the product's value (i.e., the seller is informed and the consumers are uninformed), and (ii) the seller may attempt to manipulate the consumers' beliefs by "seeding" the observational learning process with a fake purchase transaction. If the seller engages in manipulation, we assume that he is caught and removed from the market, earning zero profit, with an exogenous probability (the "detection rate") that represents the intensity of the anti-manipulation measures present in the market.[1]

The nature of our analysis is game-theoretic: Taking into account the detection rate, the seller chooses an (observable) price along with an (unobservable) manipulation strategy conditional on his type (i.e., whether his product is "good" or "bad"), and the consumers choose a purchase strategy conditional on the history of purchase observations. We assume that a bad product is worthless for the consumers, so that attention is focused on equilibria where prices are uninformative (i.e., the two seller types always pool on price) and the social learning process is activated.[2] By solving for equilibrium, we analyze the interaction between price and manipulation, and identify the impact of defensive measures on the two. By characterizing long-term learning outcomes (conditional on the players' equilibrium strategies), we investigate the associated implications for seller and consumer payoffs. Our main qualitative insights are summarized as follows.

(i) **A good seller never engages in manipulation.** This result is not readily obvious since, presumably, even a good seller may stand to benefit from interfering with the consumers' learning process. Instead, we show that the opposite is true: in any equilibrium of the game, the good seller, taking into account the bad seller's strategy, prefers to refrain from manipulation, preserving as much as possible the integrity of the social learning process. This structural result, apart from being of independent interest, significantly reduces the space of candidate equilibria, allowing us to provide a comprehensive solution to the game.

(ii) **An increase in the intensity of defensive measures does not imply a decrease in seller manipulation.** In fact, instead of deterring manipulation, we show that in all possible equilibria of the game the bad seller's manipulation strategy is locally (i.e., barring discontinuities) increasing in the detection rate. In particular, we observe that higher detection rates increase the consumers' trust in the social learning process, causing equilibrium prices to rise; at the same time, higher prices are accompanied by higher levels of manipulation by the bad seller, as the reward from successfully posing as a good seller increases.

[1] While social learning manipulation can take many forms (some examples of which are given above), most cases adhere to the following general description:

(i) There is a seller with a product of unknown value;
(ii) There is a social learning process through which consumers can learn the product's value;
(iii) There is a method by which the seller may attempt to manipulate the social learning process;
(iv) There is a form of defensive measures against such manipulations instated by a central planner.

The observational learning model we analyze is most consistent with the example of fake purchase transactions, and is chosen in this paper for its simplicity and familiarity as the "workhorse" model of social learning.

[2] Equilibria where separation on prices occurs renders the social learning process (as well as any attempts to manipulate it) redundant, and are therefore not of interest in this paper; see §3 for a related discussion.

(iii) **The absence of defensive measures is payoff-equivalent to the presence of extreme defensive measures.** Once the strategic nature of the interaction between seller and consumers is taken into account, we find that the relationship between seller manipulation and equilibrium payoffs is far from obvious. In a result that underscores this observation, we demonstrate that both the seller's and the consumers' payoffs are identical in the two extreme cases where the detection rate is either very low or very high, despite the difference in the equilibrium level of manipulation under the two (significant versus none, respectively).

(iv) **Defensive measures of intermediate intensity often yield increased payoffs for all players.** Taking the last observation one step further, we highlight a surprising phenomenon: in many cases identified in our analysis, intermediate detection rates (which admit an intermediate level of seller manipulation) yield a payoff increase for all players involved, relative to very high detection rates. In particular, we show that under an appropriately chosen detection rate, the consumers' response to seller manipulation can "nudge" the game into equilibria with lower prices and higher total welfare. Thus, far from being a bad equilibrium feature, manipulation attempts by the seller can often be leveraged to improve market efficiency.

The rest of this paper is organized as follows. In §2 we review the related literature and in §3 we present our model of social learning manipulation. In §4 we analyze the seller's and the consumers' equilibrium strategies and payoffs. In §5 we extend our analysis to allow for manipulation by the seller over multiple periods. We conclude with a discussion of our results in §6.

2. Related Literature

The literature on social learning has its origins in the seminal papers by Banerjee (1992) and Bikhchandani et al. (1992), that demonstrate how rational agents may disregard their own information in favor of mimicking the decisions of their peers. The basic model of observational social learning has since been extended in many directions and has found many applications (e.g., Acemoglu et al. 2011, Drakopoulos et al. 2013, Guarino et al. 2011, Herrera and Hörner 2013, Smith and Sørensen 2000); our model represents yet another such extension (i.e., the ability of a seller to plant fake consumer decisions in the observational learning process) and yet another application (i.e., the study of the implications of social learning manipulation) of the classic model.

A growing body of work considers the implications of various social learning processes for seller decisions. Debo et al. (2012), Debo et al. (2019), Veeraraghavan and Debo (2009), and Veeraraghavan and Debo (2011) introduce observational learning to queueing systems, investigating the implications of consumer learning from queue lengths on queue-joining behavior and system dynamics, and Kremer and Debo (2015) find experimental support for related theoretical predictions. Crapis et al. (2017), Ifrach et al. (2019), Papanastasiou and Savva (2016), Shin and Zeevi (2017) and Yu et al. (2015) analyze pricing decisions in the presence of social learning from product reviews, while Feldman et al. (2016) and Godes (2016) investigate the impact of word of mouth on product design and quality. Papanastasiou (2019b) and Hu et al. (2015) study newsvendor models in the presence of peer-to-peer social influence. The aforementioned papers assume that social learning occurs in the absence of any seller interference or deception tactics.[3] By contrast, we analyze seller decisions in a setting where the consumers' social learning process may be subject to manipulation.

While the above papers consider the implications of social learning for seller decisions, another stream of work focuses on the operations of online platforms connecting sellers to consumers.[4] Frazier et al. (2014), Papanastasiou et al. (2017) and Che and Hörner (2018) study the extent to which optimal subsidies and information disclosure can be effective in motivating self-interested consumers to "explore" different products, while Bimpikis et al. (2019) investigate the relationship between platform information disclosure and the entry/exit decisions of suppliers active on the platform. In a crowdfunding setting, Chakraborty and Swinney (2017) demonstrate how a creator's choice of crowdfunding campaign parameters might allow consumers to learn the quality of his product. Closer to our work is a recent paper by Papanastasiou (2019a) which analyzes social learning manipulation in the context of fake news. In that paper too, the authors consider the effectiveness of anti-manipulation measures in the form of penalties imposed on a social-media platform, and find that higher penalties do not necessarily lead to improved social learning outcomes. However, the mechanism underlying their result differs substantially from ours: ours is based on distortions in the strategic interaction between sellers and consumers, while theirs is based on the platform's own anti-manipulation incentives.

Our work also has strong connections to the literature on the implications of counterfeit/copycat products on seller decisions and market outcomes. Qian et al. (2014) take a signaling approach to study whether and how authentic sellers might seek to differentiate themselves from counterfeit sellers through product appearance and quality, while Pun et al. (2018) investigate the adoption of blockchain technology as an alternative signaling device. Gao et al. (2016) study copycat entry and competition dynamics when consumers derive status utility based on the average wealth of product buyers. Pun and DeYong (2017) study competition between authentic manufacturers and copycats in the presence of strategic consumer behavior, and Cho et al. (2015) highlight the difference in effective approaches for an authentic seller competing with deceptive and nondeceptive counterfeiters. In our paper, rather than copying the product attributes of a good seller, the deceptive seller attempts to "copy" the good seller's social learning process, in order to convince consumers that

[3] See Johnson and Sokol (2020) for a related discussion on AI-based collusion.

[4] See Chen et al. (2019) for an overview of work on online platforms and discussion of research opportunities.

his product is worth purchasing. In this respect, closest to our work is a recent paper by Jin et al. (2019), which analyzes how "brushing" (the practice of faking orders to improve sales ranking) affects the usefulness of ranking systems for customer search; interestingly, the authors find that in the presence of brushing, consumers may be better off under a random-ranking system.

3. Model Description

Basic Setting. We consider a model of observational learning in the mold of the classic papers by Banerjee (1992) and Bikhchandani et al. (1992). There is an infinite sequence of homogeneous consumers, each of whom must make a once-and-for-all decision regarding whether to purchase a new product. The product's value can be low or high, $V \in \{0,1\}$, and the net utility derived from purchasing a product of value V is given by $V - p$, where p is the product's price. The agents' prior belief that the product is of high value is represented by the probability $b \in (0,1)$. Apart from the information contained in the prior belief, in choosing whether to purchase the product, each consumer observes (i) the product's price, (ii) the history of the preceding consumers' purchase decisions, and (iii) the realization of a private informative signal $s \in \{0,1\}$ generated according to $P(s = v \mid V = v) = a$ for $v \in \{0,1\}$, where $a \in (0.5,1)$ is the signal's informativeness. Each consumer seeks to maximize her individual expected utility, and we normalize the utility gain from not purchasing the product to zero. If a consumer is indifferent between purchasing and not, she purchases with (an endogenously determined) probability $\mu \in [0,1]$.

The product's value is assumed to be privately known to the product seller, and if the product is of high (low) value, we say that the seller is "good" ("bad"); let $i \in \{G, B\}$ denote the seller's type.[5] At the beginning of time, the seller observes the value of his product, and chooses the product's price p_i.[6] In addition to the (observable) price, the seller also chooses whether to engage in an (unobservable) attempt to manipulate the consumers' social learning process; the details of the seller's manipulation strategy are described in the next subsection. The goal of the seller is to maximize his expected profit per period over the infinite horizon.[7]

[5] The assumption that the seller is informed is motivated by settings where "authentic" sellers coexist with sellers of counterfeit/low-quality products (see also Qian et al. (2014) and Pun et al. (2018)).

[6] Existence of equilibria in the model with manipulation will require that price is chosen from a finite set. We will assume throughout that this set is infinitely dense, so that price can be treated as a continuous variable. To simplify exposition, whenever we say that the equilibrium price is p, it should be understood that the equilibrium price is the largest element in the price set that is strictly smaller than p.

[7] From an analytical perspective, the objective of maximizing expected profit per period is significantly advantageous compared to that of maximizing total expected discounted profit. See §6 for a related discussion.

Social Learning Manipulation. We model social learning manipulation through a manipulation rate $q_i \in [0,1]$, which represents the probability that a seller of type $i \in \{G, B\}$, upon introducing his product, also plants a fake purchase to seed the observational learning process (in §5, we extend the model to allow for multi-period manipulation).[8] The cost of engaging in manipulation is captured by an exogenous detection rate $\phi \in (0,1]$, which represents the probability with which a manipulation is detected (e.g., by the platform hosting the seller or a government authority). For simplicity, we assume that if the seller is caught engaging in manipulation, he is removed from the market, earning zero profit.

Equilibrium Concept. The equilibrium concept we employ is that of Perfect Bayesian Equilibrium (PBE). A PBE requires that (i) the strategies of the two seller types are optimal given the consumers' posterior beliefs, (ii) the consumers' posterior beliefs are consistent with the sellers' strategies, and (iii) the consumers' posterior beliefs are calculated via Bayes' rule wherever possible. Based on the observable part of the seller's strategy (i.e., the product's price), there are two possible types of equilibrium outcomes: separating equilibria and pooling equilibria. In a separating equilibrium, the two seller types choose a distinct price and the consumers become informed about the product's value just by observing the product's price; in this case, social learning is redundant. By contrast, in a pooling equilibrium, both seller types choose the same price, so that the consumers cannot infer the product's value from its price; in this case, consumers engage in social learning in order to learn the product's value. To focus our discussion on equilibria where the social learning process is active (and therefore manipulation attempts are relevant), our model by design precludes the possibility of separating equilibria: in any separating equilibrium, the bad seller would extract zero profit (by merit of our assumption that a low-value product is worthless), so that there is always an incentive to pose as a good seller.[9]

We note that in all instances of our model, there exist a continuum of pooling equilibria (as is common in signaling games). While our analysis identifies all such equilibria, to draw attention to those outcomes which appear to be the most plausible in our setting, we apply the equilibrium selection approach described in Mailath et al. (1993). In particular, we apply the "undefeated" equilibrium refinement, which in our model rules out equilibria that are payoff-inferior for both seller types; moreover, in cases where more than one undefeated equilibria exist, we select the lexicographically maximum sequential equilibrium (LMSE), which effectively amounts to selecting the undefeated equilibrium that is preferred by the good seller type. We provide a more detailed discussion of our equilibrium selection along with the relevant definitions in Appendix A.

[8] Our approach to modeling manipulation is most consistent with the practice of posting fake purchases in order to inflate sales performance (see also Jin et al. (2019)).

[9] See §6 for a discussion relating to separating equilibria.

4. Analysis

We begin in §4.1 with an analysis of the benchmark case where price is endogenous but manipulation is impossible. In §4.2, we provide a comprehensive solution to the general game with social learning manipulation. We discuss the main insights obtained from our equilibrium analysis in §4.3.

4.1. Benchmark: No Manipulation

We analyze first the price-only equilibrium of our model when manipulation is impossible (i.e., for $q_G = q_B = 0$ fixed exogenously). Apart from providing a benchmark for comparison, the analysis of this simpler case illustrates several features of our analytical approach, which will be useful in the more complex case with manipulation analyzed in §4.2.

As is standard in models of observational learning, long-term outcomes are determined in our model by the emergence of "informational cascades," whereby from some point onwards, the information contained in the history of purchase observations overwhelms the information contained in any individual consumer's private signal, and all consumers simply "herd" on the purchase decision of their predecessor (see Banerjee (1992) and Bikhchandani et al. (1992)). As there are two possible consumer actions in our model, there are two types of possible cascades, namely, purchase cascades and non-purchase cascades. If a seller's product triggers a purchase (non-purchase) cascade, the seller achieves a profit per period equal to the product's price (equal to zero).

Given that our model precludes the possibility of informative prices, the expected profit per period of a seller of type $i \in \{G, B\}$ is given by $\pi_i = \Gamma_i(p, b) \cdot p$, where $\Gamma_i(\cdot, \cdot)$ is the ex ante probability that the seller's product triggers a purchase cascade and p is the product's equilibrium price. Lemma 1 characterizes the seller's purchase cascade probabilities in the benchmark case without manipulation, for an arbitrary price p.

Lemma 1. *Suppose manipulation is impossible (i.e., fix $q_G = q_B = 0$). Define $b_l < b_m < b_h$, where*

$$b_l = \frac{1}{1 + \left(\frac{a}{1-a}\right)\left(\frac{1-p}{p}\right)}, \quad b_m = \frac{1}{1 + \left(\frac{1-p}{p}\right)} = p, \quad b_h = \frac{1}{1 + \left(\frac{1-a}{a}\right)\left(\frac{1-p}{p}\right)}.$$

If the product's price is p and the consumers' prior belief is b, the purchase cascade probability of a seller of type $i \in \{G, B\}$ is $\Gamma_i(p, b)$, where

$$\Gamma_G(p,b)=\begin{cases}0 & \text{if } b\in(0,b_l),\\ \frac{a^2[\mu^2+a\mu(1-\mu)]}{1-a+a^2} & \text{if } b=b_l,\\ \frac{a^2}{1-a+a^2} & \text{if } b\in(b_l,b_m),\\ \frac{a\mu+a^2(1-\mu)}{1-a+a^2} & \text{if } b=b_m,\\ \frac{a}{1-a+a^2} & \text{if } b\in(b_m,b_h),\\ \frac{a+(1-a)^2[\mu+a\mu(1-\mu)]}{1-a+a^2} & \text{if } b=b_h,\\ 1 & \text{if } b\in(b_h,1).\end{cases}\qquad \Gamma_B(p,b)=\begin{cases}0 & \text{if } b\in(0,b_l),\\ \frac{(1-a)^2[\mu^2+(1-a)\mu(1-\mu)]}{1-a+a^2} & \text{if } b=b_l,\\ \frac{(1-a)^2}{1-a+a^2} & \text{if } b\in(b_l,b_m),\\ \frac{(1-a)\mu+(1-a)^2(1-\mu)}{1-a+a^2} & \text{if } b=b_m,\\ \frac{1-a}{1-a+a^2} & \text{if } b\in(b_m,b_h),\\ \frac{1-a+a^2[\mu+(1-a)\mu(1-\mu)]}{1-a+a^2} & \text{if } b=b_h,\\ 1 & \text{if } b\in(b_h,1).\end{cases} \tag{1}$$

All proofs are relegated to the Appendix. A useful way to analyze the consumers' learning process is to view the public belief after each observation as a discrete random walk. The purchase cascade probabilities given in Lemma 1 can then be calculated as the first-passage probabilities of the walk for appropriate belief boundaries (i.e., the two belief boundaries above and below which a purchase and non-purchase cascade is triggered, respectively). The difference between the purchase cascade probabilities of a good and bad seller is attributed to the consumers' private signals, which correctly reflect the product's underlying value (equivalently, the seller's type) with probability $a\in(0.5,1)$. Observe that the purchase cascade probabilities are step functions, owing to the discreteness of the belief updates after each purchase/non-purchase observation (see Figure 1 for an example).

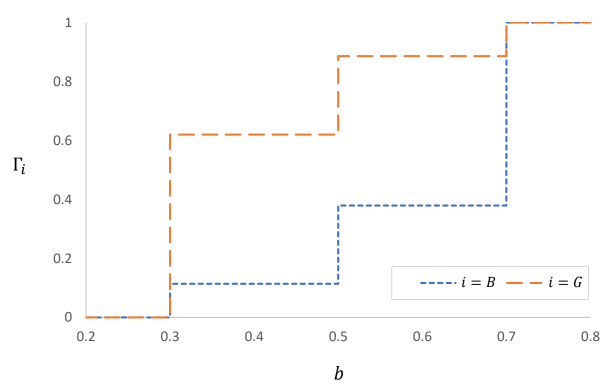

Figure 1 **Purchase cascade probability $\Gamma_i(p,b)$, as a function of the prior belief b, for a fixed price p. Parameter values: $p=0.5$ and $a=0.7$.**

It is worthwhile to pay particular attention to the special cases where $b\in\{b_l,b_m,b_h\}$, which we henceforth refer to as "indifference beliefs." In these cases, the consumers' observational learning process admits at least some observation histories at which consumers find themselves indifferent between purchasing the product and not. Whenever this is the case, consumers purchase the product with probability μ, which in turn enters the expressions for the seller's purchase cascade probabilities as given in Lemma 1. Although any $\mu \in [0,1]$ can form part of a PBE in the game without manipulation, to simplify the exposition in the remainder of this section we focus on equilibria with $\mu = 1$, without loss of generality.

To identify equilibria in the benchmark model, we define the price thresholds $p_l < p_m < p_h$, where

$$p_l = \frac{1}{1+\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}, \quad p_m = \frac{1}{1+\left(\frac{1-b}{b}\right)} = b, \quad p_h = \frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)},$$

and express the expected profit per period of each seller type as

$$\pi_G^0 = \begin{cases} p & \text{if } p \in [0,p_l], \\ \left(\frac{a}{1-a+a^2}\right)p & \text{if } p \in (p_l,p_m], \\ \left(\frac{a^2}{1-a+a^2}\right)p & \text{if } p \in (p_m,p_h], \\ 0 & \text{otherwise,} \end{cases} \qquad \pi_B^0 = \begin{cases} p & \text{if } p \in [0,p_l], \\ \left(\frac{1-a}{1-a+a^2}\right)p & \text{if } p \in (p_l,p_m], \\ \left(\frac{(1-a)^2}{1-a+a^2}\right)p & \text{if } p \in (p_m,p_h], \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

The last expressions reveal that the two sellers' expected profit functions are piecewise linear and increasing in price. We note that for a given prior belief b, any price p can be supported as a price-only pooling PBE in our model, by specifying off-equilibrium beliefs appropriately. However, it is desirable to identify and analyze the most plausible among these equilibria. To do so, we apply the equilibrium selection process described in §3. First, we note that according to the profit expressions in (2), any pooling equilibrium at a price other than $\{p_l, p_m, p_h\}$ cannot survive as an undefeated equilibrium, since both seller types can achieve a strictly higher payoff at a pooling equilibrium at slightly higher price. Next, we compare the three remaining candidate equilibria. This comparison yields at least one undefeated equilibrium; if more than one undefeated equilibria exist, we select the LMSE. The described process leads to Proposition 1.

Proposition 1. *Suppose manipulation is impossible (i.e., fix $q_G = q_B = 0$). Define $\beta_m < \beta_h$, where*

$$\beta_m = \frac{a^2-(1-a)}{2a-1} \text{ and } \beta_h = \frac{a^2-(1-a)\left(\frac{1-a}{a}\right)}{2a-1}.$$

The equilibrium price p^0 takes one of three possible values. In particular:

$$p^0 = \begin{cases} p_h & \text{for } b \in (0,\beta_m], \\ p_m & \text{for } b \in (\beta_m,\beta_h], \\ p_l & \text{for } b \in (\beta_h,1). \end{cases} \tag{3}$$

The result is illustrated in Figure 2. When the prior belief satisfies $b \in (0,\beta_m]$ the seller opts for price p_h which is accompanied by a low purchase cascade probability (see (2)). Within this range of prior beliefs, the higher the consumers' belief, the higher the price the seller can charge while maintaining the same purchase cascade probability. However, when the prior belief crosses into the region $b \in (\beta_m, \beta_h]$, the seller adopts a different approach, dropping the price to p_m so as to induce a higher purchase cascade probability. The equilibrium price then gradually increases with b until the next shift in the seller's approach occurs, whereby the seller drops the price to p_l.

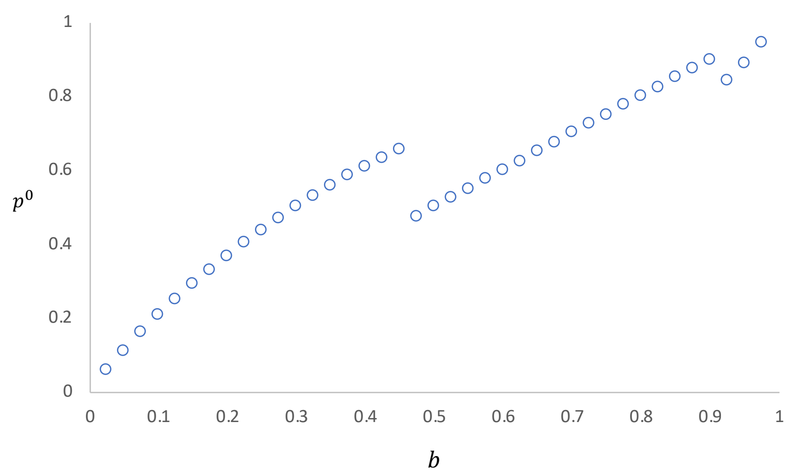

Figure 2 **Equilibrium price p^0 in the no-manipulation benchmark, as a function of the prior belief b. Parameter values: $a = 0.7$.**

To conclude this section, it will be useful to provide the functional forms for the seller's and the consumers' equilibrium payoffs. Let π (respectively, C) denote the seller's ex ante expected profit per period (the consumers' ex ante expected surplus per period); then

$$\begin{aligned} \pi &= b \cdot \pi_G(p,b) + (1-b) \cdot \pi_B(p,b) \\ &= [b \cdot \Gamma_G(p,b) + (1-b) \cdot \Gamma_B(p,b)] \cdot p, \end{aligned} \tag{4}$$

$$C = b \cdot \Gamma_G(p,b) \cdot (1-p) - (1-b) \cdot \Gamma_B(p,b) \cdot p, \tag{5}$$

where p is the equilibrium price and $\Gamma_i(p,b)$ are the seller's equilibrium purchase cascade probabilities. Observe from these expressions that the seller's ex ante payoff is increasing in the equilibrium price as well as in both purchase cascade probabilities; on the contrary, the consumers' ex ante payoff is increasing only in the good seller's purchase cascade probability, and is decreasing in the equilibrium price and in the bad seller's purchase cascade probability. Summing the expressions for profit and surplus gives the total welfare function

$$W = b \cdot \Gamma_G(p, b), \tag{6}$$

which reveals that value is generated only when a good product is successfully identified through the social learning process and adopted by the consumers. Although the presence of manipulation in the subsequent analysis will affect the equilibrium prices and purchase cascade probabilities, the functional forms of the seller's and consumers' payoffs remain as given above.

In what follows, we use the superscript "0" to denote equilibrium quantities in the no-manipulation benchmark of Proposition 1.

4.2. Equilibrium Characterization

We now return to the general model where the seller, apart from choosing the product's price, may also attempt to manipulate social learning by seeding the observational learning process with a fake purchase. In this section, we provide a complete solution to the manipulation game along with a qualitative description of our solution approach. The main insights that can be extracted from this solution are discussed subsequently in §4.3.

Expanding on the model description in §3, a PBE of the game with manipulation is summarized by (i) the the seller's strategy $\{p, q_i\}$, where p denotes the (pooling) price and q_i the manipulation strategy of a seller of type $i \in \{G, B\}$; (ii) the consumers' belief $b(h_t)$ that the product is of high value, where h_t denotes the history of purchase observations; and (iii) the purchase probability μ of an indifferent consumer. Since in any PBE of our model the bad seller always mimics the good seller on price, it is not necessary to distinguish between seller types when referring to the product's price. However, we note this does not also imply that $q_G = q_B$; since the seller's manipulation strategy is unobservable, the two seller types may adopt a different approach to manipulation.

To solve the social learning manipulation game, we leverage the fact that after the first transaction is observed by the consumers (note that this transaction may or may not be the result of manipulation), the consumers' adoption process follows the standard observational learning paradigm described in §4.1; specifically, if the product's price is p and the consumers' belief after the first observed transaction is b', then the purchase cascade probability of a seller of type i is $\Gamma_i(p, b')$, as given in Lemma 1. In turn, this structure allows us to focus our analysis on the posterior belief b', which can then be used to "map" the two sellers' payoffs to those of a standard observational learning process. With respect to the posterior belief b', there are three possible scenarios: (i) if the first observed transaction is a purchase, then $b' = b^p$, (ii) if the first observed transaction is a non-purchase, then $b' = b^n$, and (iii) if the seller attempts a manipulation and is caught, then the game is terminated with the seller achieving zero profit (equivalently, in this case we may set $b' = 0$, resulting in zero future adoption).

Even given the above structure, the problem of solving the manipulation game remains challenging, not least because the posterior beliefs b^p and b^n are endogenous to the two sellers' manipulation strategies q_G and q_B (we discuss the details of this dependence later on). The following structural result simplifies the task at hand considerably.

PROPOSITION 2. *In any PBE of the social learning manipulation game, the seller's manipulation strategy satisfies* $q_G^* = 0$ *and* $q_B^* \in [0, 1)$.

Proposition 2 suggests that in equilibrium a good seller never engages in manipulation. The result relies on two observations. The first is that engaging in manipulation is both more costly (i.e., the expected loss in profit if caught is larger) and less beneficial (i.e., the expected impact on the consumers' learning process is smaller) for a good seller. This implies that in any equilibrium, a bad seller must manipulate at a higher rate than a good seller. The second observation is that equilibria where both sellers manipulate at a strictly positive rate cannot be sustained; in particular, as we show in the proposition's proof, in any such proposed equilibrium, the good seller can always benefit by decreasing his manipulation rate (until this hits the lower bound of zero manipulation).[10]

Given that in all equilibria of our game we have $q_G^* = 0$ for the good seller, to simplify the notation we henceforth write $q := q_B$ for the bad seller's manipulation rate. Apart from establishing a result of independent interest, Proposition 2 is crucial in ensuring tractability to provide a comprehensive solution to the manipulation game; this is presented in Proposition 3 and Table 1.

PROPOSITION 3. *For any given detection rate* ϕ*, the seller's equilibrium strategy* (p^*, q^*) *in the unique undefeated LMSE takes one of the seven possible forms indexed in Table 1. In particular:*

(i) If $\phi \in (0, a^2)$*, the equilibrium takes one of the forms with index* $\{1, 2, 4, 5, 7\}$ *in Table 1.*

(ii) If $\phi \in \left[a^2, \frac{a^2}{1-a+a^2}\right)$*, the equilibrium takes one of the forms with index* $\{1, 2, 3, 5, 7\}$ *in Table 1.*

(iii) If $\phi \in \left[\frac{a^2}{1-a+a^2}, a\right)$*, the equilibrium takes one of the forms with index* $\{1, 3, 5, 7\}$ *in Table 1.*

(iv) If $\phi \in [a, 1]$*, the equilibrium takes one of the forms with index* $\{1, 3, 6\}$ *in Table 1.*

Moreover, for any given value of the detection rate ϕ*, the index of the equilibrium is non-increasing in the prior belief* b*.*

The social learning manipulation game is sufficiently complex to warrant some comments on the solution approach. As in the benchmark case of §4.1, the equilibrium characterization follows two

[10] We note that the result of Proposition 2 depends on the assumption that all consumers in our model are rational; if, by contrast, the consumer population includes a significant proportion of naive consumers (who do not adjust their beliefs according to the seller's strategy), one might expect the good seller to also engage in manipulation.

Eqm	ϕ range	p^*	q^*	μ^*
1	$(0,1]$	$p_l = \frac{1}{1+\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}$	0	1
2	$\left(0, \frac{a^2}{1-a+a^2}\right)$	$\frac{1}{1+\left(\frac{(1-\phi)}{(2a-1)(1-\mu^*)+(1-\phi)}\right)\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}$	$\frac{(2a-1)(1-\mu^*)}{(2a-1)(1-\mu^*)+(1-\phi)}$	$\frac{2-a-\sqrt{a^2+\frac{4(1-a)(1-a+a^2)\phi}{a^2}}}{2(1-a)}$
3	$[a^2,1]$	$p_m = \frac{1}{1+\left(\frac{1-b}{b}\right)}$	0	$\frac{2-a-\sqrt{(2-a)^2+\frac{4(1-a)^2(a^2-\phi)}{a^2(a-\phi)}}}{2(1-a)}$ for $\phi \in \left[a^2, \frac{a^2}{1-a+a^2}\right)$ 1 for $\phi \in \left[\frac{a^2}{1-a+a^2}, 1\right]$
4	$(0,a^2)$	$\frac{1}{1+\left(\frac{a(1-\phi)}{a-\phi(1-a)}\right)\left(\frac{1-b}{b}\right)}$	$\frac{2a-1}{(1+\phi)a-\phi}$	1 - $\frac{\sqrt{\phi}}{a}$
5	$(0,a)$	$\frac{1}{1+\left(\frac{1-\phi}{2a-\phi}\right)\left(\frac{1-b}{b}\right)}$	$\frac{2a-1}{2a-\phi}$	$\left(\frac{1-a}{a}\right)\left(\frac{\phi}{a-\phi}\right)$ for $\phi \in (0,a^2)$ $\frac{a-1+\sqrt{(1-a)^2+\frac{4(a-\phi)}{1-a}}}{2a}$ for $\phi \in [a^2,a)$
6	$[a,1]$	$p_h = \frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)}$	0	1
7	$(0,a)$	$\frac{1}{1+\left(\frac{(1-a)(1-\phi)}{a-\phi(1-a)}\right)\left(\frac{1-b}{b}\right)}$	$\frac{2a-1}{(1+\phi)a-\phi}$	$\frac{a-1+\sqrt{(1+a)^2-4\phi}}{2a}$

Table 1 List of possible equilibria in the social learning manipulation game.

steps. In the first step, we identify, essentially by construction, all possible PBE of our model. This results in a continuum of possible equilibria, belonging to seven qualitatively different equilibrium regimes. Each of the seven regimes applies to a specific range of the detection rate ϕ, with a maximum of five regimes applying for any given value of ϕ. In the second step we apply the equilibrium selection process described in §3 to identify the most plausible equilibrium of the game. Proposition 3 catalogues each of the equilibria that may emerge in the manipulation game for any given value of the detection rate ϕ, and establishes the basic structure of the equilibrium relative to the prior belief b. We now discuss each step in more detail.

To identify all possible equilibria of the game, we start with an arbitrarily chosen price p (i.e., the observable component of the seller's strategy), and ask whether and how this price can form part of a PBE. A PBE at price p requires existence of a supporting manipulation strategy (i.e., the unobservable component of the seller's strategy) which is in equilibrium with the consumers' purchase strategy. Since by Proposition 2 a good seller never engages in manipulation, it suffices to focus on the manipulation strategy of the bad seller and the consumers' response to this strategy.

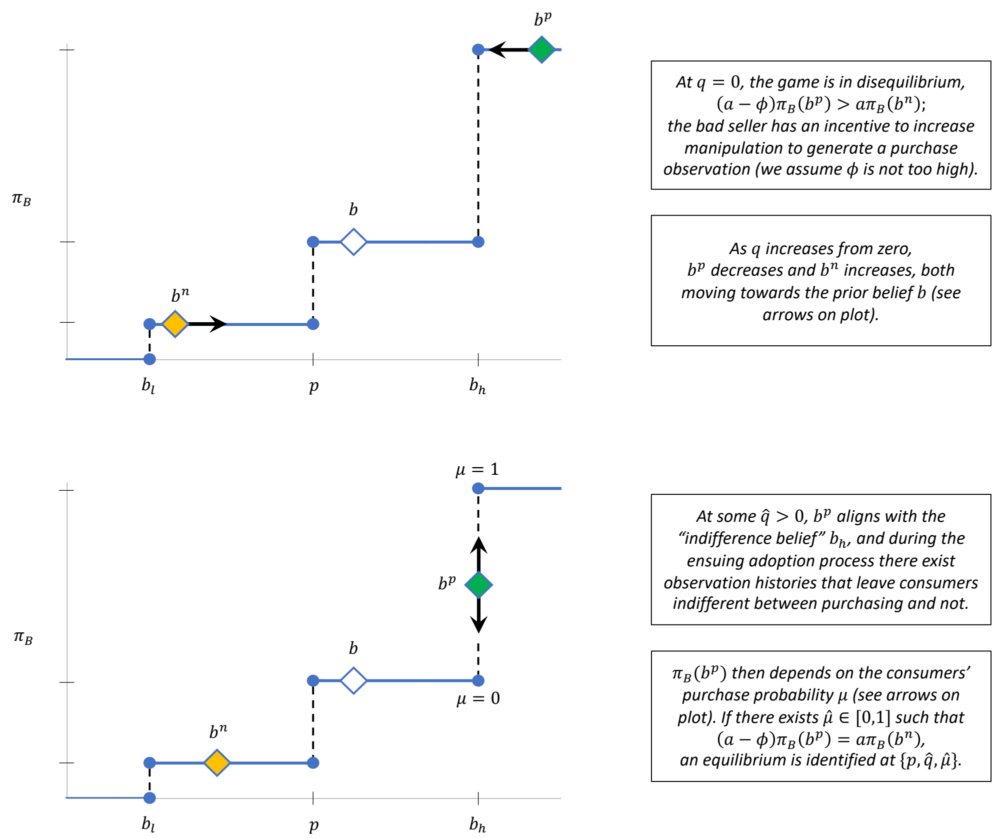

Figure 3 **Example construction of a PBE at prior belief b and price p. Note: the example assumes small ϕ.**

To help explain the process of identifying an equilibrium, we enlist the example of Figure 3. In this example we consider a price $p \in (p_l, b)$, and seek to identify a supporting manipulation strategy for the bad seller. Bypassing the technical details (see the proof of Proposition 3), it can be shown that indifference of the bad seller with respect to manipulation reduces to the equation

$$(a - \phi)\pi_B(b^p) = a\pi_B(b^n), \tag{7}$$

where $\pi_B(b^p) = \Gamma_B(p, b^p) \cdot p$ (respectively, $\pi_B(b^n) = \Gamma_B(p, b^n) \cdot p$) is the bad seller's expected profit in the event that the first observed transaction is a purchase (non-purchase). Observe first that the no-manipulation tactic of $q = 0$ in general cannot be part of an equilibrium. To see this, notice that in an equilibrium with $q=0$ the bad seller's expected profit after a purchase is observed in the first period is higher than the profit after a non-purchase is observed (i.e., $\pi_B(b^p) > \pi_B(b^n)$); therefore, assuming the detection rate is not prohibitively high (more on this later), the bad seller has an incentive to increase his manipulation activity q so as to make a purchase observation more probable. As the bad seller's manipulation rate q increases from zero, the consumers' equilibrium posterior beliefs b^p and b^n converge towards the prior belief b (see Figure 3, upper plot).[11]

As q continues to increase, b^p eventually aligns with the "indifference belief" b_h. Recall that an indifference belief implies that during the ensuing adoption process, there exist (throughout the process) observation histories which leave the next consumer indifferent between purchasing and not. Depending on the indifferent consumers' purchase probability μ, the bad seller's profit $\pi_B(b^p)$ can take a range of values (see Figure 3, lower plot). If a purchase probability μ exists such that the bad seller is indifferent between the first observed transaction being a purchase and a non-purchase (equivalently, such that $(a-\phi)\pi_B(b^p) = a\pi_B(b^n)$), an equilibrium is identified: for the given price p, there exists a supporting manipulation strategy q which is in equilibrium with the consumers' adoption strategy (which itself includes the indifferent consumer's purchase probability μ).

The above process highlights two features of equilibrium in the social learning manipulation game. The first is that for any price p, only a finite set of manipulation strategies q (i.e., those that cause either b^p or b^n to align with one of the indifference beliefs $\{b_l, b_m, b_h\}$) need be considered in identifying a PBE. The second is that equilibria with positive manipulation can be classified according to the equilibrium values of the posteriors b^p and b^n, one of which is "pinned" to one of the indifference beliefs and the other of which lies somewhere between two of the indifference beliefs. Each such combination amounts to an equilibrium "regime" and each regime can be supported across a continuous range of prices (e.g., in the example of Figure 3 the same qualitative equilibrium can be supported for any $p \in (p_l, b)$).

We next apply our equilibrium selection process to identify the most plausible equilibrium of the game. To do so, we first observe that the purchase cascade probability of each seller type is constant within each equilibrium regime. This implies that all but the highest price belonging to each regime are defeated, since both seller types can benefit by shifting play to an equilibrium belonging to the same regime but at a higher price. In total, seven candidate equilibria survive this

[11] To see why this occurs, notice that b^p must decrease in q because a purchase observation becomes more likely to have been generated by a bad seller, while b^n must increase in q because a non-purchase becomes more likely to have been generated by a good seller. In the example of Figure 3, the functional forms of the posteriors are

$$b^p(q) = \frac{1}{1+\left(\frac{1-a+q(a-\phi)}{a}\right)\left(\frac{1-b}{b}\right)} \quad \text{and} \quad b^n(q) = \frac{1}{1+(1-q)\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}.$$

process (i.e., the equilibria presented in Table 1), while, as Proposition 3 suggests, only a subset of the seven equilibria apply for any given value of the detection rate ϕ.

Finally, we compare the candidate equilibria to identify a unique undefeated equilibrium or, in the case where more than one undefeated equilibria exist, the unique undefeated LMSE. This leads to the final statement in Proposition 3. Observe that equilibria are indexed in Table 1 in order of increasing price. Therefore, the result suggests that the equilibrium in the game with manipulation exhibits a structure similar to that in Proposition 1 (see also Figure 2) with respect to the equilibrium price; that is, there exist thresholds on the prior belief b at which the equilibrium price drops abruptly and between which it increases continuously. To illustrate the equilibrium structure, we present the example of Figure 4. In this example, we have $\phi \in \left[a^2, \frac{a^2}{1-a+a^2}\right]$ so that, according to Proposition 3, the unique undefeated LMSE takes one of the forms with index $\{1,2,3,5,7\}$ in Table 1. Moreover, observe that, consistent with the final statement of Proposition 3, as the prior belief b increases from zero to one, the equilibrium shifts through forms $7 \to 5 \to 3 \to 1$.[12]

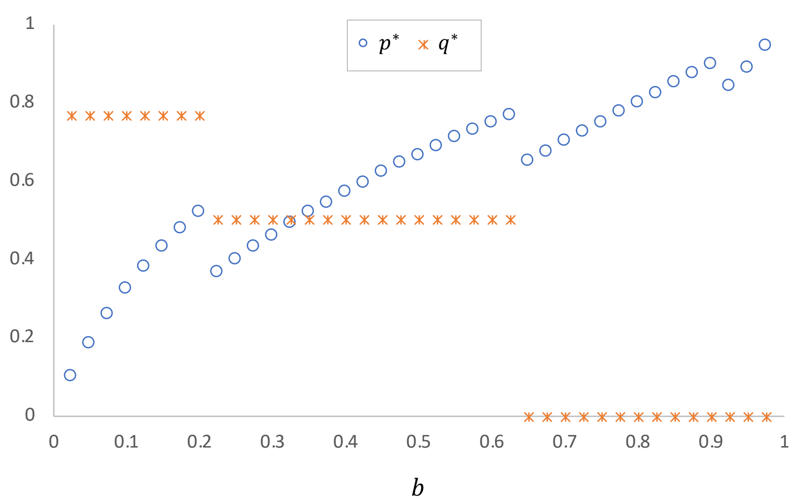

Figure 4 **Equilibrium price p^* and manipulation q^*, as a function of the prior belief b. Parameter values: $a = 0.7$, $\phi = 0.6$.**

4.3. Insights

Although Proposition 3 and Table 1 paint a complicated picture of equilibrium in the social learning manipulation game, several interesting insights can be extracted through a closer examination of the individual equilibria. We focus on insights pertaining to the interaction between the seller's

[12] Note that, in general, not all possible equilibrium forms listed in Proposition 3 necessarily emerge as the undefeated LMSE as b varies. In the example of Figure 4, the equilibrium with index 2 does not emerge as an undefeated LMSE.

and the consumers' equilibrium strategies and their respective payoffs, and how these are affected by the manipulation detection rate.

The first result of this section is an intuitive one, which has already been alluded to in the discussion of §4.2.

Corollary 1. *Suppose $\phi \in [a, 1]$. Then the equilibrium manipulation rate is $q^* = 0$ and the game reduces to the benchmark equilibrium of Proposition 1.*

Irrespective of the product's price, when the detection rate is sufficiently high (i.e., $\phi \in [a, 1]$), the chance of getting caught in the act of manipulation is sufficiently large to deter any misconduct from the bad seller, who prefers instead to mimic the good seller's pricing decision and hope for errors in the consumers' private signals to drive product adoption. Given that no manipulation occurs in equilibrium, the model then collapses to the benchmark case without manipulation discussed in §4.1, with identical equilibrium strategies and payoffs.

We note that while measures of high intensity can be effective in deterring misconduct, the cost of implementing such measures may be prohibitive for a central planner (we note that in our model a detection rate $\phi \geq a$ implies $\phi > 0.5$), so that measures of lesser intensity may warrant consideration. In addition, as we show later on, it is far from obvious that the no-manipulation equilibrium achieved by high-intensity measures is the most desirable outcome of the game (this is true irrespective of whether the central planner's objective is to maximize consumer surplus, seller profit, or a mixture of the two).

Accordingly, we next discuss insights pertaining to low-to-intermediate detection rates $\phi \in (0, a)$, starting with an interesting observation regarding the structure of the seller's equilibrium strategy.

Proposition 4. *Suppose $\phi \in (0, a)$. The equilibrium price p^* and manipulation rate q^* are both piecewise increasing in ϕ.*

The result of Proposition 4 comprises two components: (i) both price and manipulation are locally monotonic and increasing in the detection rate ϕ; and (ii) the seller's equilibrium strategy exhibits discontinuities (hence the piecewise nature of the result); we discuss each of the two in turn.

The locally increasing nature of the seller's equilibrium strategy $\{p^*, q^*\}$ follows directly from Table 1, by observing that in each of the equilibria that apply to cases of $\phi \in (0, a)$ (see Proposition 3), both price and manipulation are increasing in the detection rate ϕ. To see why this occurs, recall that in each equilibrium, one of the two posterior beliefs b^p and b^n lies on an "indifference belief," and also that each equilibrium applies to a particular range of ϕ values. As the detection rate ϕ increases within the applicable range, the equilibrium price and manipulation rate adjust so as to maintain the posterior belief b^p or b^n "pinned" to the indifference belief. In any equilibrium, this adjustment requires that the price and the manipulation rate move in the same direction. To illustrate, consider Equilibrium 7 in Table 1, which is an equilibrium involving a non-purchase posterior of $b^n = b_l = \frac{1}{1+\frac{a}{1-a}\cdot\frac{1-p}{p}}$. Maintaining this equilibrium as ϕ changes in $(0, a)$ requires

$$b^n = \frac{1}{1+(1-q)\cdot\frac{a}{1-a}\cdot\frac{1-b}{b}} = \frac{1}{1+\frac{a}{1-a}\cdot\frac{1-p}{p}},$$

revealing a positive relationship between p and q. In particular, in this example a higher price must be accompanied by a higher posterior belief following a non-purchase transaction b^n; in turn, this can only be achieved if the bad seller manipulates at a higher rate, rendering a non-purchase transaction more likely to have been generated by a good seller. Analogous relationships govern each of the equilibria of Table 1, establishing that price and manipulation are positively related.

Given that price and manipulation move in the same direction, let us consider next why the pair is increasing in the detection rate ϕ (as opposed to decreasing). As ϕ increases, the highest price under which any particular qualitative equilibrium can be supported increases, as a result of the overall positive impact of the detection policy on the consumers' beliefs: the higher the detection rate, the higher the consumers' belief that a seller who remains active in the market (i.e., who has not been caught engaging in manipulation) is a good seller. As a result, the equilibrium price, and by association the equilibrium manipulation rate, rise to higher values.

As for the second component of Proposition 4, that is, the discontinuous nature of the seller's policy, this can occur for two reasons. The first reason is consumer-driven and occurs on the boundaries of ϕ given in Proposition 3: As ϕ increases, the consumers' adoption strategy shifts so as to maintain indifference of the bad seller with respect to the manipulation strategy; however, whenever ϕ reaches one of the boundaries given in Proposition 3, the consumers' adjustment is "maxed out," so that the same qualitative equilibrium can no longer be supported. The second reason is seller-driven and occurs *between* the boundaries of Proposition 3: As ϕ increases and the consumers' adoption strategy adjusts to maintain the equilibrium, the LMSE may change, resulting in a discontinuity in the seller's equilibrium strategy.

In general, we find that the occurrence and nature of the discontinuities in the seller's equilibrium strategy as ϕ increases varies depending on the parameters of the learning model (i.e., a and b). Figure 5 presents two examples. In both examples, the seller's price and manipulation strategy are piecewise increasing in the detection rate for any $\phi < a = 0.7$ (as suggested by Proposition 4), and are constant for any $\phi \geq a$ (as suggested by Corollary 1), although the first example exhibits a much more regular pattern than the second.

As Proposition 4 and the numerical experiments of Figure 5 suggest, the impact of low-to-intermediate detection rates $\phi \in (0, a)$ on the equilibrium payoffs of the seller and the consumers

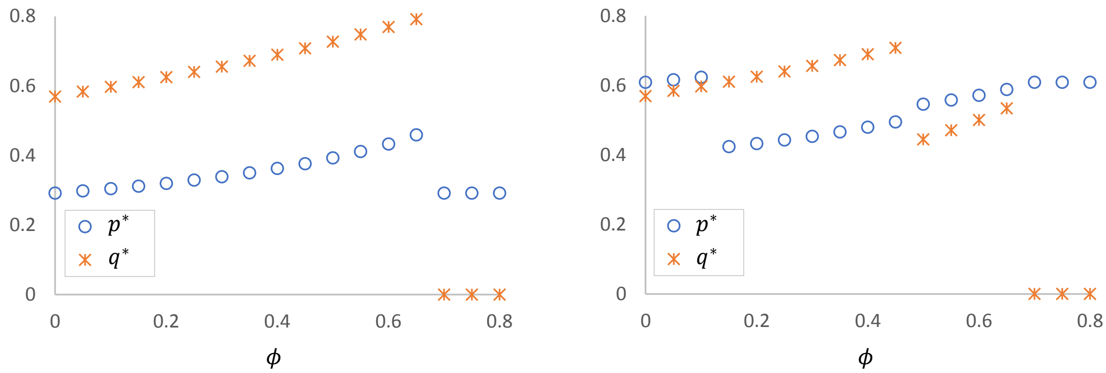

Figure 5 **Equilibrium price p^* and manipulation q^*, as a function of the detection rate ϕ. Left plot parameter values: $a = 0.7$ and $b = 0.15$. Right plot parameter values: $a = 0.7$ and $b = 0.4$.**

is in general complicated and varies depending on the parameters a and b. Rather than presenting an exhaustive analysis of payoffs for these cases, in the remainder of this section we focus attention on two significant findings.

The first of the two suggests that, to the extent that the no-manipulation benchmark equilibrium described in §4.1 represents the most desirable outcome of the game, it may not even be necessary for a central planner to bother with *any* anti-manipulation measures.[13]

PROPOSITION 5. *Suppose $\phi \to 0$.*

(i) The equilibrium manipulation rate is

$$q^* = \begin{cases} 2 - \frac{1}{a} & \text{for } b \in (0, \beta_h], \\ 0 & \text{for } b \in (\beta_h, 1), \end{cases} \tag{8}$$

and the equilibrium price is $p^* = p^0$ as given in (3).

(ii) The seller's and the consumers' expected payoffs satisfy $\pi_i^* = \pi_i^0$ for $i \in \{G, B\}$ and $C^* = C^0$.

Remarkably, although in the absence of any detection measures significant levels of seller manipulation may be observed in equilibrium, the rest of the equilibrium quantities—prices and player payoffs—are identical to those in the no-manipulation equilibrium of Proposition 1. In fact, the proof of Proposition 5 establishes an even stronger result: any PBE of the game at price p leads to seller and consumer payoffs equal to those in the no-manipulation equilibrium at price p, irrespective of the level of equilibrium manipulation.

[13] In Proposition 5, we avoid $\phi = 0$ as this leads to a continuum of payoff-equivalent equilibria; we note that the limiting equilibria discussed in the proposition are also equilibria for $\phi = 0$.

The key to the result lies in the strategic interaction between the bad seller and the consumers. Recall that the consumers' posterior belief after the first transaction is observed can be either b^p (following a purchase observation) or b^n (following a non-purchase observation); let $\pi_B(b^p)$ and $\pi_B(b^n)$ denote the seller's expected profit in these two scenarios, under some fixed price p. Now, observe that in the absence of detection, in any proposed equilibrium with $\pi_B(b^p) > \pi_B(b^n)$ (respectively, $\pi_B(b^p) < \pi_B(b^n)$), the bad seller can profitably deviate by increasing (decreasing) his manipulation rate q, so as to make a purchase observation more (less) probable. This leads to the conclusion that an equilibrium, if it exists, must be such that $\pi_B(b^p) = \pi_B(b^n)$. The requirement $\pi_B(b^p) = \pi_B(b^n)$ means that the consumers' purchase strategy must adjust in a manner such that the ensuing adoption process is "dynamics-equivalent" to a process where the first purchase/non-purchase observation is completely ignored, which in turn reduces to the benchmark case without manipulation. Thus, at any price p, the seller's expected payoff is equivalent to that in the benchmark case without manipulation; it follows that the price and equilibrium payoffs of the seller and the consumers are identical to those under the no-manipulation equilibrium of Proposition 1.

It may appear odd that the bad seller engages in manipulation, even though this brings no benefit to him in terms of expected profit. The reason is the seller's lack of ability to commit not to manipulate. To see this, consider an equilibrium where the bad seller does not manipulate; in such an equilibrium, purchase observations are more valuable than non-purchase observations, providing an incentive for the bad seller to manipulate, destroying the equilibrium. Instead, an equilibrium is reached when the bad seller can no longer benefit from increasing his manipulation rate—when there is no chance of getting caught, this can only be the case when purchase observations and non-purchase observations lead to an equivalent expected payoff, as described above. (In contrast, non-zero defensive measures disturb this balance and can have the unintended consequences described in Proposition 4, as the bad seller now also accounts for the possibility of getting caught.)

The result of Proposition 5 can be viewed as particularly positive to the extent that it suggests that no-manipulation payoffs can be achieved in the presence of manipulation, without the necessity of implementing any anti-manipulation measures. While it might appear sensible to conclude the analysis here, the next result highlights a further surprising phenomenon.

PROPOSITION 6. *Define* $\beta_l < \beta_m$, *where*

$$\beta_l = \frac{a^2-(1-a)}{a(2a+1)-1} \text{ and } \beta_m = \frac{a^2-(1-a)}{2a-1}.$$

Suppose $b \in [\beta_l, \beta_m]$. *Then there exists a detection rate* $\phi \in (a^2, a)$ *such that the seller's and the consumers' expected payoffs satisfy* $\pi_i^* > \pi_i^0$ *for* $i \in \{G, B\}$ *and* $C^* > C^0$.

Proposition 6 suggests that when the prior belief lies between the two thresholds β_l and β_m, seller manipulation can in fact be leveraged to increase the payoffs of all players, through an appropriately chosen detection rate of intermediate intensity. To illustrate the result of Proposition 6, in Figure 6 we shade the region of model parameters that satisfy the condition $b \in [\beta_l, \beta_m]$, and in Figure 7 we provide an example where we compute the increase in the payoff of the two seller types (relative to the benchmark case without manipulation) as a function of the detection rate ϕ.

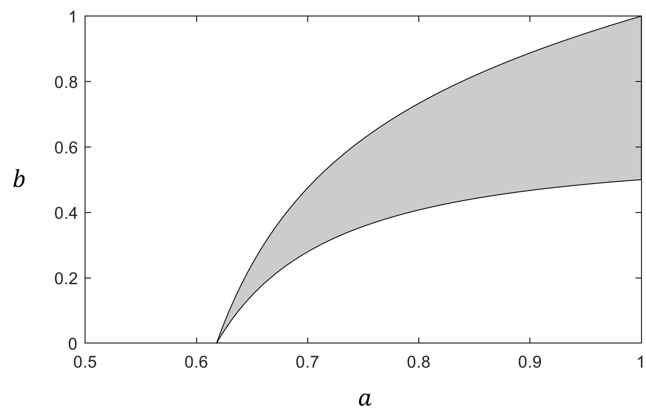

Figure 6 **Shaded region marks parameter combinations where an intermediate detection rate $\phi \in (a^2, a)$ results in higher equilibrium payoffs for all players (as compared to the no-manipulation payoffs of Proposition 1).**

To explain this phenomenon, it is instructive to discuss what happens as the detection rate ϕ crosses the boundary $\phi = a$ from above. Recall from Corollary 1 that while $\phi \geq a$, the equilibrium is identical to the benchmark equilibrium with no manipulation, and the seller's equilibrium strategy for $b \in (0, \beta_m]$ is given by (see Proposition 1)

$$p^* = \frac{1}{1 + \left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)} = p_h \quad \text{and} \quad q^* = 0.$$

As ϕ decreases and enters the region $\phi < a$, the above equilibrium "splits" into two possible equilibria, depending on the value of the prior belief b within $(0, \beta_m]$. Specifically, as we show in the proof of Proposition 6:

(i) If $b \in (0, \beta_l)$, the seller's equilibrium strategy is (see Equilibrium 7 in Table 1)

$$p^* = \frac{1}{1 + \left(\frac{(1-a)(1-\phi)}{a-\phi(1-a)}\right)\left(\frac{1-b}{b}\right)} > p_h \quad \text{and} \quad q^* = \frac{2a-1}{(1+\phi)a - \phi}.$$

(ii) If $b \in [\beta_l, \beta_m]$, the seller's equilibrium strategy is (see Equilibrium 5 in Table 1)

$$p^* = \frac{1}{1 + \left(\frac{1-\phi}{2a-\phi}\right)\left(\frac{1-b}{b}\right)} < p_h \quad \text{and} \quad q^* = \frac{2a-1}{2a-\phi}.$$

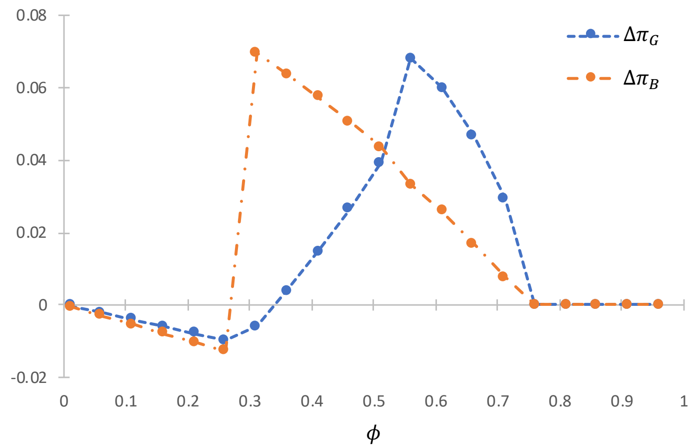

Figure 7 **Increase in the seller's equilibrium payoff ($\Delta\pi_G = \pi_G^* - \pi_G^0$ and $\Delta\pi_B = \pi_G^* - \pi_B^0$) relative to the benchmark case ($\pi_G^0 = 0.52$ and $\pi_B^0 = 0.06$), as a function of the detection rate ϕ. Parameter values: $a = 0.75$ and $b = 0.5$.**

Both of the two possible equilibria at $\phi < a$ involve positive manipulation in equilibrium; however, crucially, the first is one of higher price and lower purchase cascade probabilities for both seller types, while the second is one of lower price and higher purchase cascade probabilities (as compared to the equilibrium at $\phi \geq a$).[14] Proposition 6 then establishes that the latter, low-price equilibrium results in higher payoffs for both the seller and the consumers.

In particular, we note first that according to expression (6), it is straightforward to deduce that the higher purchase cascade probability for a good seller results in an increase in total welfare. The less obvious and more interesting aspect of the result is that this increase in total welfare is enjoyed by both seller types as well as the consumers, all of whom are better off under the new equilibrium. For the consumers (see expression (5)), the higher probability of identifying a good product coupled with the decrease in price results in significant surplus gains, and even though some of these gains are attenuated by an increase in the purchase cascade probability of a bad seller, the net effect remains positive. Moreover, for each of the two seller types (see expression (4)), the purchase cascade probability increases to an extent which more than compensates for the reduction in the equilibrium price (though the increase in purchase cascade probability is proportionately higher for the good type), leaving both seller types with a higher equilibrium payoff.

[14] This can be deduced from the values of p^* and μ^* in Table 1 in the limit as ϕ approaches a from below.

5. Model Extension: Multi-Period Manipulation

In this section, we extend our model and analysis to the case of multi-period manipulation. In particular, we now assume that the seller can manipulate over T periods, for any finite $T \geq 1$.[15] The history of observed transactions in period t is denoted by $h_t \in \mathcal{H}_t$, where $\mathcal{H}_t$ is the set of possible histories. At the beginning of the selling horizon, the seller chooses the product's price p.[16] Moreover, at the beginning of each period, the seller (taking into account the history of observed transactions) chooses a manipulation strategy denoted by $q_{i,t}(h_t) \in [0,1]$ for $i \in \{G, B\}$. The probability that the seller is caught (conditional on engaging in manipulation) is $\phi \in (0,1]$ in each period.

We begin the analysis of the multi-period manipulation model with a result that establishes the robustness of Proposition 2 in our main analysis.

PROPOSITION 7. *In any PBE of the multi-period manipulation game, the seller's manipulation strategy satisfies $q^*_{G,t} = 0$ and $q^*_{B,t} \in [0,1)$.*

Remarkably, the proof of Proposition 7 not only establishes the existence of an equilibrium in the dynamic multi-period manipulation game, it also shows that in any such equilibrium, the good seller refrains from manipulation at all times and at all possible observation histories. Although the result is significantly more involved than its single-period manipulation counterpart from a technical perspective, the intuition according to which the good seller does not engage in manipulation carries through. In particular, manipulation for the good seller is both less beneficial and more costly, so that in equilibrium his manipulation rate must be lower than that of the bad seller. At the same time, equilibria where both seller types manipulate at a positive rate cannot be supported; as a result, in any PBE of the game, only the bad seller's manipulation strategy can be positive. Moreover, as is the case in our main model, the only way to guarantee that the bad seller never engages in manipulation is to enforce defensive measures of signifiant intensity.

PROPOSITION 8. *Suppose $\phi \in [a, 1]$. Then the bad seller's equilibrium manipulation rate satisfies $q^*_{B,t} = 0$ and the game reduces to the benchmark equilibrium of Proposition 1.*

Intuitively, irrespective of whether the seller can conduct manipulation over multiple periods, if the intensity of defensive measures is sufficiently high, the seller will never engage in manipulation. As a result, the model again reduces to the benchmark model considered in §4.1.

We next consider whether the insight of Proposition 5 regarding the payoffs of the seller and the consumers in the absence of defensive measures holds in the case of multi-period manipulation.

[15] Our main model corresponds to the case $T = 1$. In the model with multi-period manipulation, a period is defined by the occurrence of an observed transaction.

[16] It is straightforward to verify that in any equilibrium both seller types choose the same price.

PROPOSITION 9. *Suppose $\phi \to 0$. Then the seller's and the consumers' expected payoffs satisfy $\pi_i^* = \pi_i^0$ for $i \in \{G, B\}$ and $C^* = C^0$.*

Proposition 9 thus extends the singe-period result of Proposition 5 to the case of multi-period manipulation. It suggests that, in the absence of defensive measures, the seller's ability to engage in manipulation does not affect either the seller's or the consumers' equilibrium payoffs—even though manipulation attempts can occur over multiple periods, the consumers' beliefs and actions adapt accordingly so that these attempts ultimately result in payoffs that are no different to those of the benchmark model without manipulation.

We move next to the case of defensive measures of intermediate intensity $\phi \in (0, a)$. Proposition 10 establishes the basic structural properties of the seller's dynamic manipulation strategy.

PROPOSITION 10. *Suppose $\phi \in (0, a)$. In any PBE of the multi-period manipulation game:*

(i) If $b_t < p^*$, then $q_{B,t}^*(b_t) > 0$.

(ii) If $b_t \geq p^*$, then there exist thresholds $\underline{\rho}_t(b_t) \leq \overline{\rho}_t(b_t) < a$ such that $q_{B,t}^*(b_t) > 0$ for any $\phi < \underline{\rho}_t(b_t)$ and $q_{B,t}^*(b_t) = 0$ for any $\phi \geq \overline{\rho}_t(b_t)$. Moreover, if $b_t > p^*$, then $\underline{\rho}_t(b_t) = \overline{\rho}_t(b_t)$.

Observe first that the structure suggested by Proposition 10 is consistent with the result of Proposition 3, in that the seller's manipulation strategy in each period depends on the public belief regarding the product's value (as well as on the intensity of defensive measures present in the system). Next, note that Proposition 10 also provides insight into how the seller's manipulation tactics evolve over time, depending on the prevailing public belief b_t. When the belief b_t is lower than the product's price p^*, the seller manipulates at a positive rate irrespective of the intensity of defensive measures ϕ. Here, in the absence of manipulation, one further non-purchase observation would result in a non-purchase cascade, providing the bad seller with a strong incentive to manipulate. When the belief is higher than the product's price, the incentive to manipulate is weaker, in that manipulation erodes the positive impact of a purchase observation on the public belief, while at the same time even if a non-purchase occurs in the current period, the seller still has an opportunity to avoid a non-purchase cascade by manipulating in the next period. As a result, when $b_t \geq p^*$ the bad seller manipulates only provided the probability of getting caught is not too high.

6. Discussion

This paper develops a theoretical model of social learning manipulation. Using this model, we derive insights regarding the strategic interaction between sellers and consumers and how this is impacted by the implementation of defensive measures aimed at deterring misconduct. Our analysis yields several findings that may be of interest to e-commerce practitioners, managers of online platforms, and government officials. The prevailing conclusion from our results is that the phenomenon of social learning manipulation is one with complex implications which may have been previously misunderstood.

Common sense might suggest that more stringent anti-manipulation measures should be more effective at deterring misconduct, causing a decrease in seller manipulation and an increase in consumer surplus. However, our analysis suggests that the opposite is often the case: anti-manipulation measures often result in increased prices, which are accompanied by increased levels of seller manipulation and decreased consumer surplus (Proposition 4). Our analysis also casts doubt on the notion that seller manipulation is symptomatic of an inefficient system. In fact, we find that when (a) consumers are reasonably good at identifying low-quality products, and (b) the prior belief that any given product is of high quality is relatively high (Figure 6), the most desirable market outcome for both sellers and consumers is one that involves positive levels of seller manipulation (Proposition 6). Moreover, even in cases where seller manipulation is not beneficial, we find that investment in anti-manipulation measures may not be warranted, as such measures do not necessarily imply an improvement in either seller profit or consumer surplus (Proposition 5).

A significant simplification in our analysis is the assumption that the bad seller's product is worthless for the consumers, which in turn rules out the possibility of price-separating equilibria and allows us to focus on equilibria where the social learning process is active. Assuming alternatively that a bad product has some positive value for the consumers, an interesting question for future work is how the presence of seller manipulation affects the parameter regions in which price-pooling versus price-separating equilibria prevail. The current paper provides the foundation for identifying and analyzing pooling equilibria, but the payoff implications for sellers and consumers of a "switch" between separating and pooling equilibria are as of yet unclear and worth investigating.

Allowing for dynamic pricing presents yet another opportunity for future research. As a first investigation into the effects of social learning manipulation, we have focused on the simpler case of fixed pricing. While there is a growing literature on dynamic pricing in the presence of social learning (e.g., Papanastasiou and Savva 2016), this literature has not yet addressed the issue of seller manipulation. Based on the current analysis, we expect the interplay between dynamic pricing and social learning manipulation to involve significant technical challenges, but may yield additional valuable insights. Another simplification of our analysis related to model dynamics is our assumption that the firm maximizes expected profit per period over an infinite horizon. This significantly enhances model tractability, but it also implicitly assumes that the firm is sufficiently forward-looking so that the short-term impact of engaging in manipulation is unimportant. Future work might consider the more general objective function of maximizing total expected discounted profit, which may provide more nuanced results relating to the firm's discounting of future profits.

Finally, it is important to recognize that social learning manipulation in practice can take many forms, including fake purchases, fake reviews and fake product endorsements, among others (see also §1). The model presented in this paper is one of action-based observational learning and is therefore most closely related to the case of manipulation through fake purchases. Although we expect several of our model insights to extend to other forms of manipulation as well, further work is necessary in order to improve our understanding of the more general phenomenon of social learning manipulation and its implications. Moreover, while this paper has focused on the basic operational lever of performing content inspections to curb seller manipulation, future work may consider alternative levers such as platform design and information disclosure, which may represent less direct, but nevertheless effective, ways of decreasing sellers' incentives to manipulate the social learning process.

Appendix

A. Equilibrium Selection: Undefeated Criterion and LMSE

Our process for selecting a unique equilibrium in the social learning manipulation game consists of two steps and is based on the approach described in Mailath et al. (1993). First, among all PBE of the game, we identify those equilibria which are "undefeated." If a unique equilibrium emerges, the process is concluded. If there are more than one undefeated equilibria, we proceed to the second step, where we compare the undefeated equilibria and select the "lexicographically maximum sequential equilibrium" (LMSE). We now provide more details on each step.

The *undefeated criterion* is developed in Mailath et al. (1993) and is motivated therein by a discussion of the shortcomings associated with other approaches to selecting equilibria in signaling games, including the intuitive criterion proposed in Cho and Kreps (1987) (see also the related discussion in Schmidt et al. (2015)). The main premise of the criterion is that it does not make sense to adjust beliefs at some out-of-equilibrium information sets without simultaneously adjusting beliefs at other information sets, including those on the equilibrium path. Moreover, once all necessary adjustments are made, the game should be at an equilibrium. This leads to the following test for identifying whether an equilibrium can be "defeated." Consider a proposed equilibrium and a message which is not sent by any sender type in this equilibrium. Suppose there is an alternative equilibrium in which the given message is sent by some non-empty set of sender types, and that this set is exactly the set of types who prefer the alternative equilibrium. Then, the receiver's beliefs at that message in the original equilibrium should be consistent with this set. If the beliefs are not consistent, the proposed equilibrium is defeated by the alternative equilibrium.

To make the discussion more concrete, we now provide the formal definition of undefeated equilibrium as it applies in our model setup. For ease of exposition, in this section we use the following notation. Let $\nu(i, h_t) = \{p(i), q(i, h_t)\}$ denote the seller's pricing-and-manipulation strategy and $\mu(p, h_t)$ the consumers' adoption strategy. Moreover, let $b(i)$ denote the consumers' prior belief that the seller is of type i and $\xi(i \mid p)$

the consumers' posterior belief that the seller is of type i conditional on observing price p at the beginning of the selling horizon.

DEFINITION 1. The PBE $\sigma = \{\nu, \mu, \xi\}$ *defeats* the PBE $\sigma' = \{\nu', \mu', \xi'\}$ if there exists a price $\hat{p}$ that satisfies the following conditions:

(i) $p'(i) \neq \hat{p}$ for all i, but $K = \{i : p(i) = \hat{p}\} \neq \emptyset$;

(ii) $\pi_i(\sigma) \geq \pi_i(\sigma')$ for all $i \in K$, and $\pi_i(\sigma) > \pi_i(\sigma')$ for at least some $i \in K$; and

(iii) there exists $i \in K$ such that $\xi'(i \mid \hat{p}) \neq \frac{b(i)\gamma(i)}{\sum_{j\in\{G,B\}} b(j)\gamma(j)}$ for any $\gamma : \{G, B\} \mapsto [0,1]$ satisfying

  (a) $i \in K$ and $\pi_i(\sigma) > \pi_i(\sigma')$ implies $\gamma(i) = 1$, and

  (b) $i \notin K$ implies $\gamma(i) = 0$.

Condition (i) requires that the price $\hat{p}$ does not feature in equilibrium σ', but is part of another equilibrium of the game σ. Condition (ii) checks that all seller types which use price $\hat{p}$ in equilibrium σ are weakly better off in equilibrium σ as compared to equilibrium σ', with at least one of these types being strictly better off.[17] Condition (iii) establishes that any off-equilibrium beliefs that can be used to support equilibrium σ' are unreasonable given the existence and properties of equilibrium σ. When the above conditions hold simultaneously, equilibrium σ is said to defeat equilibrium σ'. A PBE is "undefeated" if no other equilibrium exists that defeats it.

Combining Definition 1 with the fact that our game admits only price-pooling equilibria, any PBE in our model which can be Pareto-improved upon can be defeated. Furthermore, while existence of an undefeated equilibrium is guaranteed, uniqueness is not (i.e., for some combinations of our model parameters, there exist more than one equilibria where improving the payoff of one seller type by shifting play to an alternative price-pooling equilibrium would cause a reduction in the other type's payoff). In cases where there is no unique undefeated equilibrium, we select among the undefeated equilibria using the idea of lexicographically maximum sequential equilibrium (LMSE).

DEFINITION 2. The PBE $\sigma = \{\nu, \mu, \xi\}$ *lexicographically dominates* the PBE $\sigma' = \{\nu', \mu', \xi'\}$ if either (i) $\pi_G(\sigma) > \pi_G(\sigma')$, or (ii) $\pi_G(\sigma) = \pi_G(\sigma')$ and $\pi_B(\sigma) > \pi_B(\sigma')$.

That is, the LMSE is the equilibrium in which the good seller's payoff is at its maximum or, conditional on the good seller's payoff being at its maximum, the bad seller's payoff is at its maximum. Selecting the LMSE in cases where the game admits multiple undefeated equilibria recognizes that in our model it is the bad seller who wishes to pose as a good seller (and not the other way around), so that it is reasonable to resolve these case on the basis of a belief structure that is consistent with such a situation.

B. Proofs

Proof of Lemma 1

Since the seller knows his type, a good (bad) seller knows that good private signals are received by the consumers with probability a (respectively, $1-a$), and bad signals with probability $1-a$ (respectively, a)

[17] A common misconception regarding the undefeated refinement is that it is equivalent to eliminating equilibria by performing a Pareto-comparison. However, observe that by Condition (ii), it is in general possible for equilibrium σ to defeat equilibrium σ' even if it results in a decrease in the payoff of some player types, namely, those that do not use price $\hat{p}$ in equilibrium σ.

in each period. As long as consumers are able to infer their predecessors' private signals from their actions, it can be verified that bad signals "cancel out" with good signals, so that what matters for the public belief is the accumulated difference between the number of good and bad signals at any point in time. Below, we conduct our analysis from the good seller's perspective, with the understanding that the corresponding result for the bad seller can be obtained by simply replacing a with $1-a$.

When the prior is b, and the consumer receives a good signal, her expected utility from purchase is

$$\frac{ab}{ab+(1-a)(1-b)}-p.$$

If this expected utility is negative, the consumer will not purchase. However, following this non-purchase observation, the public belief remains b (because the consumer's private signal cannot be inferred from her action). Thus, all subsequent consumers will choose not to purchase, resulting in a non-purchase cascade. It follows that any prior

$$b<b_l:=\frac{1}{1+\left(\frac{a}{1-a}\right)\left(\frac{1-p}{p}\right)} \iff p>p_h:=\frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)}$$

results in a non-purchase cascade with certainty, i.e., $\Gamma_G=0$. Following a similar logic, any prior

$$b>b_h=\frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-p}{p}\right)} \iff p<p_l:=\frac{1}{1+\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}$$

results in a purchase cascade, i.e., $\Gamma_G=1$.

Now consider the case $b\in[b_l,b_h]$, in which case the actions of the first few agents are (at least partially) informative with respect to their private signals. After n net good signals have accumulated, the public belief is

$$b'=\frac{a^nb}{a^nb+(1-a)^n(1-b)}.$$

For the system to enter a purchase cascade, we require that the next consumer in line chooses to purchase even if her private signal is negative. That is, we require

$$\frac{a^{n-1}b}{a^{n-1}b+(1-a)^{n-1}(1-b)}>p.$$

Rearranging, we have

$$U=\min\left\{n\in\mathbb{Z}:n>1-\frac{\ln\left(\frac{1-p}{p}\frac{b}{1-b}\right)}{\ln\left(\frac{a}{1-a}\right)}\right\},$$

where U is the minimum number of net good signals that triggers a purchase cascade. Similarly, it can be shown that a non-purchase cascade is triggered when the number of net bad signals reaches

$$D=\min\left\{n\in\mathbb{Z}:n>1+\frac{\ln\left(\frac{1-p}{p}\frac{b}{1-b}\right)}{\ln\left(\frac{a}{1-a}\right)}\right\}.$$

Observe that from the seller's perspective the consumer adoption process is a simple random walk, and the probability of a purchase cascade occurring is equivalent to the probability that the difference between the number of good signals and the number of bad signals reaches U before it reaches $-D$ (i.e., this is a version of the "gambler's ruin"). The probability of a purchase cascade for the good seller is therefore

$$\Gamma_G = \frac{1-\left(\frac{1-a}{a}\right)^D}{1-\left(\frac{1-a}{a}\right)^{U+D}}.$$

Now, observe that in the expressions for U and D above, we have $\ln\left(\frac{a}{1-a}\right) > 0$. Therefore, if $p = b_m < b < b_h$, we have $U = 1$ and $D = 2$ and the probability of a purchase cascade for a good seller is given by

$$\Gamma_G = \frac{1-(\frac{1-a}{a})^2}{1-(\frac{1-a}{a})^3} = \frac{a}{1-a+a^2}.$$

Alternatively, if $b_l < b < b_m = p$, we have $U = 2$ and $D = 1$ and the probability of a purchase cascade for a good seller is given by

$$\Gamma_G = \frac{1-(\frac{1-a}{a})}{1-(\frac{1-a}{a})^3} = \frac{a^2}{1-a+a^2}.$$

Next, consider the three boundary cases $b \in \{b_l, b_m, b_h\}$, also referred to in the main text as "indifference beliefs." First, suppose $b = b_m = p$. In this case, when the consumers receive a good signal followed by a bad signal, the posterior returns to p, and the consumer will purchase with a probability $\mu \in [0,1]$. For a given μ, when there are two consecutive purchases the posterior belief is

$$\frac{ba[a+(1-a)\mu]}{ba[a+(1-a)\mu]+(1-b)(1-a)[(1-a)+a\mu]} > b = p.$$

Hence, a purchase cascade is formed. When there are two consecutive non-purchases, the posterior belief is

$$\frac{b(1-a)[(1-a)+a(1-\mu)]}{b(1-a)[(1-a)+a(1-\mu)]+(1-b)a[a+(1-a)(1-\mu)]} < b = p.$$

Hence, a non-purchase cascade is formed. When there is one purchase and one non-purchase, the posterior belief is $b = p$. It follows that the good seller's purchase cascade probability is

$$\Gamma_G = a[a+(1-a)\mu]+[a(1-a)(1-\mu)+(1-a)a\mu]\Gamma_G \iff \Gamma_G = \frac{a\mu+a^2(1-\mu)}{1-a+a^2}.$$

Similarly, in the case of $b = b_h$, the good seller's purchase cascade probability is

$$\Gamma_G = a+(1-a)\mu+(1-a)(1-\mu)\frac{a^2(1-\mu)+a\mu}{1-a+a^2} = \frac{a+(1-a)^2[\mu+a\mu(1-\mu)]}{1-a+a^2},$$

while in the case of $b = b_l$, the good seller's purchase cascade probability is

$$\Gamma_G = a\mu\frac{a^2(1-\mu)+a\mu}{1-a+a^2} = \frac{a^2[\mu^2+a\mu(1-\mu)]}{1-a+a^2}.$$

In summary, the full characterization of Γ_G, the purchase cascade probability of a good seller, is

$$\Gamma_G = \begin{cases} 0 & \text{if } b \in (0, b_l), \\ \frac{a^2[\mu^2+a\mu(1-\mu)]}{1-a+a^2} & \text{if } b = b_l, \\ \frac{a^2}{1-a+a^2} & \text{if } b \in (b_l, b_m), \\ \frac{a\mu+a^2(1-\mu)}{1-a+a^2} & \text{if } b = b_m, \\ \frac{a}{1-a+a^2} & \text{if } b \in (b_m, b_h), \\ \frac{a+(1-a)^2[\mu+a\mu(1-\mu)]}{1-a+a^2} & \text{if } b = b_h, \\ 1 & \text{if } b \in (b_h, 1). \end{cases}$$

By symmetry, replacing a with $(1-a)$ in the above expressions gives the purchase cascade probability of a bad seller,

$$\Gamma_B = \begin{cases} 0 & \text{if } b \in (0, b_l), \\ \frac{(1-a)^2[\mu^2+(1-a)\mu(1-\mu)]}{1-a+a^2} & \text{if } b = b_l, \\ \frac{(1-a)^2}{1-a+a^2} & \text{if } b \in (b_l, b_m), \\ \frac{(1-a)\mu+(1-a)^2(1-\mu)}{1-a+a^2} & \text{if } b = b_m, \\ \frac{1-a}{1-a+a^2} & \text{if } b \in (b_m, b_h), \\ \frac{1-a+a^2[\mu+(1-a)\mu(1-\mu)]}{1-a+a^2} & \text{if } b = b_h, \\ 1 & \text{if } b \in (b_h, 1). \end{cases}$$

□

Proof of Proposition 1

Note from the expressions in (2) that the payoffs of both seller types are piecewise linear and strictly increasing in p. Consider the region of prices $p \in [0, p_l]$. Observe that any PBE at price $p \in [0, p_l)$ results in strictly lower payoff for both seller types as compared to the PBE at price p_l. Thus, any equilibrium at price $p \in [0, p_l)$ is defeated. Repeating the same argument for the price regions $p \in (p_l, p_m]$ and $p \in (p_m, p_h]$ reveals that the game admits at most three undefeated equilibria, namely, $p^* \in \{p_l, p_m, p_h\}$ (existence of at least one undefeated equilibrium is straightforward to verify).

Next, we identify the LMSE by comparing the good seller's payoff in each of the three candidate equilibria (note that by Definition 2 in Appendix A, an LMSE is guaranteed to exist and to be undefeated). Using the expressions in (2), the good seller's equilibrium payoff in each equilibrium is given by

$$\left\{ p_l, \frac{ap_m}{1-a+a^2}, \frac{a^2 p_h}{1-a+a^2} \right\}.$$

Comparing these profits, we identify the equilibrium price which maximizes the good seller's profit for any given value of b. Specifically, comparing p_l and $\frac{ap_m}{1-a+a^2}$, we have

$$p_l > \frac{ap_m}{1-a+a^2} \iff \frac{1}{1+\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)} > \frac{ab}{1-a+a^2} \iff b > \frac{a^2-(1-a)\left(\frac{1-a}{a}\right)}{2a-1} = \beta_h.$$

Comparing $\frac{ap_m}{1-a+a^2}$ and $\frac{a^2 p_h}{1-a+a^2}$, we have

$$\frac{ap_m}{1-a+a^2} > \frac{a^2 p_h}{1-a+a^2} \iff \frac{ab}{1-a+a^2} > \left(\frac{a^2}{1-a+a^2}\right)\left(\frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)}\right) \iff b > \frac{a^2-(1-a)}{2a-1} = \beta_m.$$

Note that $\beta_h > \beta_m$. It follows from the above comparison that the LMSE price is

$$p^0 = \begin{cases} p_h & \text{if } b \in (0, \beta_m], \\ p_m & \text{if } b \in (\beta_m, \beta_h], \\ p_l & \text{if } b \in (\beta_h, 1). \end{cases}$$

□

Proof of Proposition 2

Consider a proposed equilibrium (p, q_G, q_B, μ) with consumer beliefs after the first transaction b^p and b^n for a purchase and non-purchase, respectively. Note first that $q_G = q_B = 1$ cannot be part of an equilibrium, since in this case the consumers would simply disregard the first transaction; that is, there is no benefit to manipulation, while there is a positive probability ϕ of getting caught which provides an incentive for both sellers to decrease manipulation—this destroys the equilibrium.

Excluding, then, the possibility of $q_G = q_B = 1$, the beliefs below are derived from our model setup and straightforward applications of Bayes' rule. In particular, there are two relevant cases in our analysis:

- Case (i). If the equilibrium is such that a real consumer entering a system with no transaction history would purchase if and only if she receives a positive private signal, then the public belief after the first transaction is

$$b^p = \frac{1}{1 + \left(\frac{1-a+(a-\phi)q_B}{a+(1-a-\phi)q_G}\right)\left(\frac{1-b}{b}\right)},$$
$$b^n = \frac{1}{1 + \left(\frac{1-q_B}{1-q_G}\right)\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}.$$

- Case (ii). If the equilibrium is such that a real consumer entering an empty system would purchase if she receives a positive private signal, and would be indifferent and thus randomize (i.e., purchase with probability μ) if she receives a negative private signal, then the public belief after the first transaction is

$$b^p = \frac{1}{1 + \left(\frac{1-a+a\mu+(a(1-\mu)-\phi)q_B}{a+(1-a)\mu+((1-a)(1-\mu)-\phi)q_G}\right)\left(\frac{1-b}{b}\right)},$$
$$b^n = \frac{1}{1 + \left(\frac{1-q_B}{1-q_G}\right)\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}.$$

Note that b^n has identical expressions in Case (i) and Case (ii). Moreover, the b^p expression in Case (i) can be viewed as a special case of those in Case (ii) by setting $\mu \equiv 0$. For ease of exposition, below, we conduct our analysis using the more general expression of Case (ii), with the understanding that Case (i) is included by fixing $\mu \equiv 0$.

The ex ante payoffs of the two sellers at the proposed equilibrium (p, q_G, q_B, μ) are

$$\pi_G = (q_G(1-\phi) + (1-q_G)(a + (1-a)\mu))\pi_G(b^p) + (1-q_G)(1-a)(1-\mu)\pi_G(b^n),$$
$$\pi_B = (q_B(1-\phi) + (1-q_B)(1-a+a\mu))\pi_B(b^p) + (1-q_B)a(1-\mu)\pi_B(b^n),$$

where $\pi_i(b^p)$ and $\pi_i(b^n)$ are the payoff functions of seller type i after first transaction occurs. Note that when the price is p, the payoff functions are $\pi_i(\cdot) = \Gamma_i(p, \cdot) \cdot p$, where $\Gamma_i(p, \cdot)$ are given in Lemma 1.

From Lemma 1, $\pi_i(b^p)$ and $\pi_i(b^n)$ are piecewise constant with respect to b^p and b^n for a given p, which implies that they are also piecewise constant in q_i, for $i \in \{G, B\}$. Thus, at any differentiable point, we have

$$\frac{d\pi_G}{dq_G} = (1-a)(1-\mu)[\pi_G(b^p) - \pi_G(b^n)] - \phi\pi_G(b^p),$$
$$\frac{d\pi_B}{dq_B} = a(1-\mu)[\pi_B(b^p) - \pi_B(b^n)] - \phi\pi_B(b^p).$$

Consider the bad seller first. Noting that $q_B \in [0,1]$, we first point out that in any equilibrium, we have either (i) $q_B = 0$ and $\frac{d\pi_B}{dq_B} < 0$, or (ii) $q_B \in (0,1)$ and $\frac{d\pi_B}{dq_B} = 0$. To rule out the possibility $q_B = 1$ and $\frac{d\pi_B}{dq_B} > 0$, observe that if $q_B = 1$, then $b^n = 1$; this implies that $\frac{d\pi_B}{dq_B} < 0$, so that the bad seller can profitably deviate by decreasing q_B. Thus, $q_B = 1$ cannot be part of an equilibrium. Thus, in the remainder we restrict attention to proposed equilibria satisfying $\frac{d\pi_B}{dq_B} \leq 0$.

Consider now the good seller. Suppose first that in the proposed equilibrium we have $b^p \leq b^n$ so that $\pi_G(b^p) \leq \pi_G(b^n)$. It follows that $\frac{d\pi_G}{dq_G} < 0$, so that if such an equilibrium exists it must be that $q_G = 0$, which proves the result.

Suppose alternatively that $b^p > b^n$, which implies $\pi_G(b^p) \geq \pi_G(b^n)$ and $\pi_B(b^p) \geq \pi_B(b^n)$. Observe that since $\pi_G(b^p) \geq \pi_B(b^p)$, if we also have $a[\pi_B(b^p) - \pi_B(b^n)] > (1-a)[\pi_G(b^p) - \pi_G(b^n)]$, then we have $\frac{d\pi_G}{dq_G} < \frac{d\pi_B}{dq_B}$. Since in any such equilibrium it must be that $\frac{d\pi_B}{dq_B} \leq 0$, it follows that $\frac{d\pi_G}{dq_G} < 0$, which in turn implies $q_G = 0$, proving the result.

According to the above discussion, it then suffices to verify whether the condition $a[\pi_B(b^p) - \pi_B(b^n)] > (1-a)[\pi_G(b^p) - \pi_G(b^n)]$ holds under the proposed equilibrium. To do so, we note that for general ϕ, a necessary condition at equilibrium is that either b^p or b^n are equal to one of the indifference posteriors $\{b_l, p, b_h\}$ (we provide a more extensive discussion on this point in the proof Proposition 3); we therefore consider all proposed equilibria that satisfy this necessary condition.

We note first that it can be verified using Lemma 1 that $a[\pi_B(b^p) - \pi_B(b^n)] > (1-a)[\pi_G(b^p) - \pi_G(b^n)]$ holds in all proposed equilibria satisfying one of the following:

- $b^p \in (b_h, 1)$ and $b^n \in \{b_l, p, b_h\}$
- $b^p = b_h$ and $b^n \in \{(0, b_l), (b_l, p), (p, b_h)\}$
- $b^p \in (p, b_h)$ and $b^n \in \{b_l, p\}$
- $b^p = p$ and $b^n \in \{(b_l, p)\}$

That is, in all of the above cases we have $\frac{d\pi_G}{dq_G} < \frac{d\pi_B}{dq_B}$, which proves the result.

However, there is one potential equilibrium regime where the sufficient condition $a(\pi_B(b^p) - \pi_B(b^n)) > (1-a)(\pi_G(b^p) - \pi_G(b^n))$ does not hold: $b^p \in (b_l, p)$ and $b^n = b_l$. For this regime, we compare $\frac{d\pi_G}{dq_G}$ and $\frac{d\pi_B}{dq_B}$ directly in order to establish that $\frac{d\pi_G}{dq_G} < \frac{d\pi_B}{dq_B}$. We have

$$
\begin{aligned}
\frac{d\pi_G}{dq_G} &= (1-a)[\pi_G(b^p) - \pi_G(b^n)] - \phi\pi_G(b^p) \\
&= (1-a)\left(\frac{a^2}{1-a+a^2} - \frac{a\mu(\mu a + (1-\mu)a^2)}{1-a+a^2}\right) - \phi\left(\frac{a^2}{1-a+a^2}\right), \\
\frac{d\pi_B}{dq_B} &= a[\pi_B(b^p) - \pi_B(b^n)] - \phi\pi_B(b^p) \\
&= a\left(\frac{(1-a)^2}{1-a+a^2} - \frac{(1-a)\mu(\mu(1-a) + (1-\mu)(1-a)^2)}{1-a+a^2}\right) - \phi\frac{(1-a)^2}{1-a+a^2}.
\end{aligned}
$$

Then,

$$
\begin{aligned}
\frac{d\pi_B}{dq_B} - \frac{d\pi_G}{dq_G} &\propto a(1-a)\left(((1-a)-a) - \mu(\mu((1-a)-a) + (1-\mu)((1-a)^2 - a^2))\right) - \phi\left((1-a)^2 - a^2\right) \\
&= \left(a(1-a)(1-\mu) - \phi\right)(1-2a).
\end{aligned}
$$

The last expression is positive provided $a(1-a)(1-\mu)<\phi$, or $a(1-a)\mu+\phi>a(1-a)$. If $\mu=1$, this condition holds trivially and thus $\frac{d\pi_G}{dq_G}<\frac{d\pi_B}{dq_B}$, which proves the result. If $\mu<1$, recall that a necessary condition for this case to be an equilibrium is $\frac{d\pi_B}{dq_B}\leq 0$, or

$$a\left(\frac{(1-a)^2}{1-a+a^2}-\frac{(1-a)\mu(\mu(1-a)+(1-\mu)(1-a)^2)}{1-a+a^2}\right)-\phi\frac{(1-a)^2}{1-a+a^2}\leq 0.$$

This implies $a^2\mu^2+a(1-a)\mu-a+\phi\geq 0$, or equivalently

$$a(1-a)\mu+\phi\geq a-a^2\mu^2=a(1-a\mu)>a(1-a),$$

for any $\mu<1$. Therefore, for any $\mu<1$ we have $\frac{d\pi_G}{dq_G}<\frac{d\pi_B}{dq_B}$, which proves the result. □

Proof of Proposition 3

Let P^p (P^n) denote the probability that a purchase (non-purchase) observation occurs in the first period. At a given price p, the bad seller's expected profit at manipulation q is

$$\pi_B(q)=[P^p\Gamma_B(p,b^p)+P^n\Gamma_B(p,b^n)]\cdot p,$$

where $\Gamma_B(\cdot,\cdot)$ is given in Lemma 1. Consider first the probabilities P^p and P^n. If the bad seller manipulates at rate q, then these probabilities are

$$P^p=q(1-\phi)+(1-q)Y \text{ and } P^n=(1-q)(1-Y),$$

where Y here denotes the probability that a real consumer arriving to a system with no observation history chooses to purchase (resulting in an authentic purchase observation). Without loss of generality, we focus here on the case where a real first consumer purchases if and only if she receives a positive private signal (we make adjustments as needed for special cases below). Accordingly, since the probability that a consumer receives a good private signal when the product is bad is $(1-a)$, we have $Y=(1-a)$, so that

$$P^p=1-a+(a-\phi)q \text{ and } P^n=a(1-q).$$

Consider next the posterior beliefs b^p and b^n. If the bad seller manipulates at rate q and the detection rate is ϕ, then the posterior beliefs are

$$b^p(q)=\frac{1}{1+\left(\frac{1-a+(a-\phi)q}{a}\right)\left(\frac{1-b}{b}\right)},$$
$$b^n(q)=\frac{1}{1+(1-q)\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}.$$

Now, observe that $\Gamma_B(p,b)$ is a step function of b for fixed p (see Lemma 1). The derivative of $\pi_B(q)$ at any differentiable point is

$$\frac{d\pi_B}{dq}=[(a-\phi)\Gamma_B(p,b^p(q))-a\Gamma_B(p,b^n(q))]\cdot p. \tag{9}$$

We make the following observations. First, if either (i) $\phi\geq a$ or (ii) $\phi>\frac{a^2}{1-a+a^2}$ and $b>p$, then $\frac{d\pi_B}{dq}<0$ so that in equilibrium we have $q^*=0$ and the game reverts to the benchmark case of Proposition 1; observe that for these cases, the equilibria 1, 3, 6 in Table 1 correspond to the equilibrium identified in Proposition 1. Second, if $0 < \phi < a$, then a necessary condition at equilibrium is that $\Gamma_B(p, b^p(q)) > \Gamma_B(p, b^n(q))$ which by Lemma 1 implies $b^p(q) > b^n(q)$; that is, a purchase observation is accompanied by a higher posterior belief than a non-purchase observation. Third, an equilibrium where $b^p(q) > b^n(q)$ and neither $b^p(q)$ nor $b^n(q)$ lie on an indifference belief $\{b_l, b_m, b_h\}$ cannot exist for general ϕ, since the probabilities Γ_i are constants.

The above observations prompt a search for equilibria at general ϕ in which either $b^p(q)$ or $b^n(q)$ lie on an indifference belief, so that the consumers' purchase probability when indifferent μ can be chosen and adjusted appropriately to sustain an equilibrium. In the following, we start from an arbitrary price and search for an equilibrium by identifying a supporting bad-seller manipulation strategy which is in equilibrium with the consumers' purchase strategy (this includes the randomization probability of an indifferent consumer).

Equilibria at prices $p > b$.

When $p > b$, it can be shown that there are three possible equilibrium regimes, specified by the following posterior beliefs: (i) $b^p = b_m = p$, or (ii) $b^n = b_l$ or (iii) $b^n = b_m = p$.

(i) Equilibria with $b^p = p$. For $b^p = p$, the manipulation strategy is $q = \frac{a}{a-\phi}\left[\left(\frac{1-p}{p}\right)\left(\frac{b}{1-b}\right) - \frac{1-a}{a}\right]$, and it can be shown that a first real consumer will purchase if and only if she receives a positive private signal provided $p_4 < p < p_h$, where $p_4 := \frac{1}{1+\left(\frac{a(1-\phi)}{a-\phi(1-a)}\right)\left(\frac{1-b}{b}\right)}$. Moreover, for such a q to form part of an equilibrium we require that either $b^n < b_l$ or $b_l < b^n < p$ (otherwise, the general equilibrium condition $b^p > b^n$ is violated), and that there exists μ such that the bad seller is indifferent between the first observed transaction being a purchase or a non-purchase.

Suppose first that $b^n < b_l$. We have

$$\pi_B = [(a-\phi)\Gamma_B(p,p)] \cdot p,$$
$$\frac{d\pi_B}{dq} = (a-\phi)\left(\frac{\mu(1-a)+(1-\mu)(1-a)^2}{1-a+a^2}\right) \cdot p > 0,$$

so that for any μ, an equilibrium at the proposed value of q cannot exist.

Suppose next that $b^n \in (b_l, p)$. We have

$$\pi_B = [(a-\phi)\Gamma_B(p,p) - a\Gamma_B(p,b^n)] \cdot p,$$
$$\frac{d\pi_B}{dq} = \left[(a-\phi)\left(\frac{\mu(1-a)+(1-\mu)(1-a)^2}{1-a+a^2}\right) - a\frac{(1-a)^2}{1-a+a^2}\right] \cdot p.$$

Indifference of the bad seller (i.e., $\frac{\pi_B}{dq} = 0$) requires

$$\mu = \left(\frac{1-a}{a}\right)\left(\frac{\phi}{a-\phi}\right),$$

so that such an equilibrium exists if and only if the above quantity lies in $[0,1]$, or equivalently, if and only if $\phi < a^2$. To complete the characterization of this equilibrium regime, we identify the range of price such that $b^n \in (b_l, p)$. It can be shown that $b^n < p$ provided $p > p_4 = \frac{1}{1+\left(\frac{a(1-\phi)}{a-\phi(1-a)}\right)\left(\frac{1-b}{b}\right)}$, and that $b^n > b_l$ provided $p < p_5 := \frac{1}{1+\left(\frac{1-\phi}{2a-\phi}\right)\left(\frac{1-b}{b}\right)}$. In summary, for $0 < \phi < a^2$ there exist a continuum of equilibria described as follows: $p_4 < p^* < p_5$, $q^* = \frac{a}{a-\phi}\left[\left(\frac{1-p^*}{p^*}\right)\left(\frac{b}{1-b}\right) - \frac{1-a}{a}\right]$, $\mu^* = \left(\frac{1-a}{a}\right)\left(\frac{\phi}{a-\phi}\right)$. Observe that since μ^* is independent of p^*, all such equilibria result in an identical purchase cascade probability for both seller types; therefore, all but Equilibrium 5 in Table 1 are eliminated according to the selection process described §3.

(ii) Equilibria with $b^n = b_l$. For $b^n = b_l$, the manipulation strategy is $q = 1 - \left(\frac{1-p}{p}\right)\left(\frac{b}{1-b}\right)$ and it can be shown that a first real consumer purchases if and only if she receives a positive private signal. Moreover, for such a q to form part of an equilibrium we require that either $b^p \in (p, b_h)$ or $b^p \in (b_l, p)$ (otherwise the general condition $b^p > b^n$ is violated; note that it is impossible in this case to have $b^p > b_h$).

Suppose first that $b^p \in (p, b_h)$. We have

$$\frac{d\pi_B}{dq} = \left[(a-\phi)\frac{(1-a)}{1-a+a^2} - a(1-a)\mu\left(\frac{\mu(1-a)+(1-\mu)(1-a)^2}{1-a+a^2}\right)\right]\cdot p.$$

Indifference of the bad seller requires

$$\mu\left(\mu(1-a)+(1-\mu)(1-a)^2\right) = \frac{a-\phi}{a}. \tag{10}$$

The last equation admits a unique solution in $[0,1]$ provided $a^2 \le \phi \le a$. To complete the characterization of this equilibrium regime, we identify the range of prices such that $b^p \in (p, b_h)$. It can be shown that $b^p < b_h$ provided $p > b$, and $b^p > p$ provided $p < p_5 = \frac{1}{1+\left(\frac{1-\phi}{2a-\phi}\right)\left(\frac{1-b}{b}\right)}$. In summary, for $a^2 \le \phi \le a$ there exist a continuum of equilibria described as follows: $b < p^* < p_5$, $q^* = 1 - \left(\frac{1-p^*}{p^*}\right)\left(\frac{b}{1-b}\right)$, $\mu^* = \frac{a-1+\sqrt{(1-a)^2+\frac{4(a-\phi)}{1-a}}}{2a}$. Observe that since μ^* is independent of p^*, all such equilibria result in an identical purchase cascade probability for both seller types; therefore, all but Equilibrium 5 in Table 1 are eliminated according to the selection process described §3.

Suppose next that $b^p \in (b_l, p)$. We have

$$\frac{d\pi_B}{dq} = \left[(a-\phi)\frac{(1-a)^2}{1-a+a^2} - a(1-a)\mu\left(\frac{\mu(1-a)+(1-\mu)(1-a)^2}{1-a+a^2}\right)\right]\cdot p.$$

Indifference of the bad seller requires

$$\mu(1-a+a\mu) = \frac{(a-\phi)}{a}.$$

The last equation admits a unique solution in $[0,1]$ provided $0 < \phi \le a$. To complete the characterization of this equilibrium regime, we identify the range of prices such that $b^p \in (b_l, p)$. It can be shown that $b^p < p$ provided $p > p_5$ and $b^p > b_l$ provided $p < p_7 := \frac{1}{1+\left(\frac{(1-a)(1-\phi)}{a-\phi(1-a)}\right)\left(\frac{1-b}{b}\right)}$. In summary, for $0 < \phi \le a$ there exist a continuum of equilibria described as follows: $p_5 < p^* < p_7$, $q^* = 1 - \left(\frac{1-p^*}{p^*}\right)\left(\frac{b}{1-b}\right)$, $\mu^* = \frac{a-1+\sqrt{(1+a)^2-4\phi}}{2a}$. Observe that since μ^* is independent of p^*, all such equilibria result in an identical purchase cascade probability for both seller types; therefore, all but Equilibrium 7 in Table 1 are eliminated according to the selection process described §3.

(iii) Equilibria with $b^n = p$. For $b^n = p$, the manipulation strategy is $q = 1 - \left(\frac{1-p}{p}\right)\left(\frac{b}{1-b}\right)\left(\frac{1-a}{a}\right)$, and a first real consumer will purchase if she receives a positive signal, and will randomize (purchase with probability μ) if she receives a bad signal. For such a q to form part of an equilibrium we require that $b^p \in (p, b_h)$ (otherwise the general condition $b^p > b^n$ is violated; note that it is impossible to have $b^p > b_h$).

Suppose that $b^p \in (p, b_h)$. We have

$$\pi_B(q) = \left[[q(1-\phi)+(1-q)(1-a+a\mu)]\left(\frac{1-a}{1-a+a^2}\right) + [(1-q)a(1-\mu)]\left(\frac{\mu(1-a)+(1-\mu)(1-a^2)}{1-a+a^2}\right)\right]\cdot p,$$

$$\frac{d\pi_B}{dq} = \left[a(1-\mu)\left(\left(\frac{1-a}{1-a+a^2}\right) - \left(\frac{\mu(1-a)+(1-\mu)(1-a^2)}{1-a+a^2}\right)\right) - \phi\left(\frac{1-a}{1-a+a^2}\right)\right]\cdot p.$$

Indifference of the bad seller requires

$$a^2(1-\mu)^2 = \phi.$$

The last equation admits a solution in $[0,1]$ provided $\phi \le a^2$. To complete the characterization of this equilibrium regime, we identify the conditions on price such that $b^p \in (p, b_h)$. It can be shown that $b^p < b_h$ provided $p > p_2(\mu) := \frac{1}{1+\left[\frac{(1-\phi)}{(2a-1)(1-\mu)+(1-\phi)}\right]\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}$ and that $b^p > p$ provided $p < p_4$. In summary, for $0 < \phi \le a^2$ there exist a continuum of equilibria described as follows: $p_2(\mu^*) < p^* < p_4$, $q^* = 1 - \left(\frac{1-p^*}{p^*}\right)\left(\frac{b}{1-b}\right)\left(\frac{1-a}{a}\right)$, $\mu^* = 1 - \frac{\sqrt{\phi}}{a}$. Observe that since μ^* is independent of p^*, all such equilibria result in an identical purchase cascade probability for both seller types; therefore, all but Equilibrium 4 in Table 1 are eliminated according to the selection process described §3.

Equilibria at prices $p \le b$.

When $b \ge p$ it can be shown that there are three possible equilibrium regimes, specified by the following posterior beliefs: (i) $b^p = b_h$, or (ii) $b^n = p$.

(i) Equilibria with $b^p = b_h$. For $b^p = b_h$, the manipulation strategy is $q = \frac{1-a}{a-\phi}\left[\left(\frac{1-p}{p}\right)\left(\frac{b}{1-b}\right) - 1\right]$, and a first real consumer will purchase if and only if she receives a positive private signal, provided $p > p_1 := \frac{1}{1+\left(\frac{1-\phi}{2a-\phi}\right)\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}$. Moreover, for such a q to form part of an equilibrium we require that either $b^n \in (b_l, p)$ or $b^n \in (p, b_h)$ (otherwise the general condition $b^p > b^n$ is violated; note that it is impossible to have $b^n < b_l$).

Suppose first that $b^n \in (b_l, p)$. We have

$$\frac{d\pi_B}{dq} = \left[(a-\phi)\left((1-a+a\mu) + a(1-\mu)\frac{\mu(1-a)+(1-\mu)(1-a)^2}{1-a+a^2}\right) - a\left(\frac{(1-a)^2}{1-a+a^2}\right)\right]\cdot p.$$

Indifference of the bad seller requires

$$(1-a+a\mu)\left(1-a\mu+a^2\mu\right) = \frac{a(1-a)^2}{(a-\phi)}.$$

The last equation admits a unique solution in $[0,1]$ provided $a^2 < \phi < \frac{a^2}{1-a+a^2}$. To complete the characterization of this equilibrium regime, we identify the range of prices such that $b^n \in (b_l, p)$. It can be shown that $b^n < p$ provided $p > p_l$ and $b^n > b_l$ provided $p < b$. In summary, for $a^2 < \phi < \frac{a^2}{1-a+a^2}$ there exist a continuum of equilibria described as follows: $p_l < p^* < b$, $q^* = \frac{1-a}{a-\phi}\left[\left(\frac{1-p^*}{p^*}\right)\left(\frac{b}{1-b}\right) - 1\right]$, $\mu^* = \frac{2-a-\sqrt{(2-a)^2+\frac{4(1-a)^2(a^2-\phi)}{a^2(a-\phi)}}}{2(1-a)}$. Observe that since μ^* is independent of p^*, all such equilibria result in an identical purchase cascade probability for both seller types; therefore, all but Equilibrium 3 in Table 1 are eliminated according to the selection process described §3.

Suppose next that $b^n \in (p, b_h)$. Indifference of the bad seller requires

$$(1-a+a\mu)\left(1-a\mu+a^2\mu\right) = \frac{a(1-a)}{(a-\phi)}.$$

The last equation admits a unique solution $\mu \in [0,1]$ provided $\phi < \frac{a^3}{1-a+a^2}$. To complete the characterization of this equilibrium regime, we identify the conditions on price such that $b^n \in (p, b_h)$. It can be shown that $b^n > p$ provided $p < p_l$ and $b^n < b_h$ provided $p > \frac{1}{1+\left(\frac{(1-\phi)a}{a-\phi(1-a)}\right)\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}$. However, $p < p_l$ violates the assumption that a first real consumer will purchase if and only if she receives a positive private signal, so that such an equilibrium cannot exist.

(ii) Equilibria with $b^n = p$ For $b^n = p$, the manipulation strategy is $q = 1 - \left(\frac{1-p}{p}\right)\left(\frac{b}{1-b}\right)\left(\frac{1-a}{a}\right)$ and a first real consumer will purchase if she receives a positive signal, and a first real consumer will purchase if she receives a positive signal, and will randomize (purchase with probability μ) if she receives a bad signal. For such a q to form part of an equilibrium we require that $b^p \in (p, b_h)$ or $b^p > b_h$ (otherwise the general condition $b^p > b^n$ is violated).

Suppose first that $b^p \in (p, b_h)$. We have

$$\pi_B(q) = \left[[q(1-\phi) + (1-q)(1-a+a\mu)]\left(\frac{1-a}{1-a+a^2}\right) + (1-q)a(1-\mu)\left(\frac{\mu(1-a)+(1-\mu)(1-a^2)}{1-a+a^2}\right)\right]\cdot p,$$
$$\frac{d\pi_B}{dq} = \left[a(1-\mu)\left(\left(\frac{1-a}{1-a+a^2}\right) - \left(\frac{\mu(1-a)+(1-\mu)(1-a^2)}{1-a+a^2}\right)\right) - \phi\left(\frac{1-a}{1-a+a^2}\right)\right]\cdot p.$$

Indifference of the bad seller requires

$$a^2(1-\mu)^2 = \phi.$$

The last equation admits a solution $\mu \in [0,1]$ provided $\phi \le a^2$. To complete the characterization of this equilibrium regime, we identify the range of prices such that $b^p \in (p, b_h)$. It can be shown that $b^p < b_h$ provided $p > p_2(\mu)$ and that $b^p > p$ provided $p < p_4$. In summary, for $0 < \phi \le a^2$ there exist a continuum of equilibria described as follows: $p_2(\mu^*) < p^* < p_4$, $q^* = 1 - \left(\frac{1-p^*}{p^*}\right)\left(\frac{b}{1-b}\right)\left(\frac{1-a}{a}\right)$, $\mu^* = 1 - \frac{\sqrt{\phi}}{a}$. Observe that since μ^* is independent of p^*, all such equilibria result in an identical purchase cascade probability for both seller types; therefore, all but Equilibrium 4 in Table 1 are eliminated according to the selection process described §3.

Suppose next that $b^p > b_h$. We have

$$\pi_B(q) = \left[[q(1-\phi) + (1-q)(1-a+a\mu)] + (1-q)a(1-\mu)\left(\frac{\mu(1-a)+(1-\mu)(1-a^2)}{1-a+a^2}\right)\right]\cdot p,$$
$$\frac{d\pi_B}{dq} = \left[a(1-\mu)\left(1 - \frac{\mu(1-a)+(1-\mu)(1-a^2)}{1-a+a^2}\right) - \phi\right]\cdot p.$$

Indifference of the bad seller requires

$$a^2(1-\mu)\left(\frac{1-(1-a)\mu}{1-a+a^2}\right) = \phi.$$

The last equation admits a unique solution $\mu \in [0,1]$ provided $\phi < \frac{a^2}{1-a+a^2}$. To complete the characterization of this equilibrium regime, we identify the range of prices such that $b^p > b_h$. It can be shown that $b^p > b_h$ provided $p < p_2(\mu)$. In summary, for $0 < \phi \le a^2$ there exist a continuum of equilibria described as follows: $p^* < p_2(\mu^*)$, $q^* = 1 - \left(\frac{1-p^*}{p^*}\right)\left(\frac{b}{1-b}\right)\left(\frac{1-a}{a}\right)$, $\mu^* = \frac{2-a-\sqrt{a^2+\frac{4(1-a)(1-a+a^2)\phi}{a^2}}}{2(1-a)}$. Observe that since μ^* is independent of p^*, all such equilibria result in an identical purchase cascade probability for both seller types; therefore, all but Equilibrium 2 in Table 1 are eliminated according to the selection process described §3.

As for the final statement in Proposition 3 (i.e., that the equilibrium index is non-increasing in b), note first that the equilibrium selection process described in §3 (see also Appendix A) selects the undefeated equilibrium that maximizes the good seller's profit. Let $p_j^*(b)$ denote the equilibrium price with index j in Table 1. Observe that each candidate equilbrium price can be written $p_j^*(b) = \frac{1}{1+z_j(a,\phi)\left(\frac{1-b}{b}\right)}$, where $z_j(a,\phi)$ is decreasing in index j for given values of a and ϕ. It follows that $p_j^*(b)$ is increasing in the index j. Also observe that the good seller's profit in the equilibrium with index j is given by $\pi_{G,j}(b) = p_j^*(b)\Gamma^*_{G,j}$, where

$\Gamma^*_{G,j}$ is independent of b. Now let $\Delta_{j,k}(b) = \pi_{G,j}(b) - \pi_{G,k}(b)$ be the difference between the good seller's profits under two equilibria with indices $j < k$. It is straightforward to show that

$$\frac{d\Delta_{j,k}(b)}{db} = \frac{\left(p_j^*(b)\right)^2 z_j(a,\phi)}{b^2}\Gamma^*_{G,j} - \frac{\left(p_k^*(b)\right)^2 z_k(a,\phi)}{b^2}\Gamma^*_{G,k} = \frac{p_j^*(b) z_j(a,\phi)}{b^2}\pi_{G,j}(b) - \frac{p_k^*(b) z_k(a,\phi)}{b^2}\pi_{G,k}(b),$$

where it can also be verified that $p_j^*(b)z_j(a,\phi) \geq p_k^*(b)z_k(a,\phi)$ for any $b \in (0,1)$. Therefore, for any given indices $j<k$, if $\Delta_{j,k}(b) = \pi_{G,j}(b) - \pi_{G,k}(b) \geq 0$ for some $b \in (0,1)$, then $\frac{d\Delta_{j,k}(b)}{db} \geq 0$. This implies that if the equilibrium with index k is dominated by the equilibrium with index j for some $b \in (0,1)$, then the same is true for any $\hat{b} > b$. We conclude that the index of the LMSE is non-increasing in b. □

Proof of Corollary 1

This follows from the case $\phi \in [a,1]$ of Proposition 3.

Proof of Proposition 4

Consider each of the possible equilibria given in Table 1. For equilibria 1, 3, and 6, we have $\frac{dp^*}{d\phi} = 0$ and $\frac{dq^*}{d\phi} = 0$. For equilibrium 2, we have $\mu^* = \frac{2-a-\sqrt{a^2+\frac{4(1-a)(1-a+a^2)\phi}{a^2}}}{2(1-a)}$, which is decreasing in ϕ. Thus, $\frac{1-\mu^*}{1-\phi}$ is increasing in ϕ, which implies $\frac{1-\phi}{(2a-1)(1-\mu^*)+(1-\phi)}$ is decreasing in ϕ and $\frac{(2a-1)(1-\mu^*)}{(2a-1)(1-\mu^*)+(1-\phi)}$ is increasing in ϕ; therefore, $\frac{dp^*}{d\phi} > 0$ and $\frac{dq^*}{d\phi} > 0$. For equilibrium 4, $\frac{1-\phi}{a-\phi(1-a)}$ is decreasing in ϕ and $\frac{2a-1}{a-(1-a)\phi}$ is increasing in ϕ; therefore, $\frac{dp^*}{d\phi} > 0$ and $\frac{dq^*}{d\phi} > 0$. For equilibrium 5, $\frac{1-\phi}{a-\phi}$ is decreasing in ϕ and $\frac{2a-1}{2a-\phi}$ is increasing in ϕ; therefore, $\frac{dp^*}{d\phi} > 0$ and $\frac{dq^*}{d\phi} > 0$. For equilibrium 7, $\frac{1-\phi}{a-\phi(1-a)}$ is decreasing in ϕ and $\frac{2a-1}{a-(1-a)\phi}$ is increasing in ϕ; therefore, $\frac{dp^*}{d\phi} > 0$ and $\frac{dq^*}{d\phi} > 0$. □

Proof of Proposition 5

From Proposition 3, we have that if $\phi \in (0, a^2)$, the possible equilibria are 1, 2, 4, 5 and 7 (see Table 1). For $\phi \to 0$, the following hold. First, under equilibrium 2, we have $p^* \to p_l$ and $q^* \to 0$. Thus, equilibrium 2 collapses to equilibrium 1 for $\phi \to 0$. Second, under equilibrium 7, $p^* \to p_h$, $q^* \to 2 - \frac{1}{a}$, and $\mu^* \to 1$. It can be verified that the purchase cascade probabilities are $\Gamma_G \to \frac{a^2}{1-a+a^2}$ and $\Gamma_B \to \frac{(1-a)^2}{1-a+a^2}$, which are identical to those under equilibrium 6 in the case $\phi \in [a,1)$. Thus, for $\phi \to 0$, equilibrium 7 is payoff-equivalent to equilibrium 6 in the case $\phi \in [a,1)$. Third, under equilibrium 4, $p^* \to p_m$, $q^* \to 2 - \frac{1}{a}$, and $\mu^* \to 1$. It can be verified that the purchase cascade probabilities are $\Gamma_G \to \frac{a}{1-a+a^2}$ and $\Gamma_B \to \frac{1-a}{1-a+a^2}$, which are identical to those under equilibrium 3 in the case $\phi \in [a,1)$. Thus, for $\phi \to 0$, equilibrium 4 is payoff-equivalent to equilibrium 3 in the case $\phi \in [a,1)$. Finally, under equilibrium 5, $p^* \to \frac{1}{1+\left(\frac{1}{2a}\right)\left(\frac{1-b}{b}\right)} < p_h$, $q^* \to 2 - \frac{1}{a}$, and $\mu^* \to 0$. It can be verified that the purchase cascade probabilities are $\Gamma_G \to \frac{a^2}{1-a+a^2}$ and $\Gamma_B \to \frac{(1-a)^2}{1-a+a^2}$, which are identical to those under equilibrium 7. Since the price under equilibrium 7 approaches p_h for $\phi \to 0$, it follows that for $\phi \to 0$ equilibrium 5 is defeated by equilibrium 7.

Therefore, for $\phi \to 0$, there are three possible undefeated equilibria at the prices corresponding to equilibria 1, 3, 6 in the case $\phi \in [a,1)$. Combining with the above discussion, the equilibrium manipulation rates under the three equilibria are

$$q^* = \begin{cases} 2 - \frac{1}{a} & \text{in the equilibria with prices } p^* \in \{p_m, p_h\}, \\ 0 & \text{in the equilibrium with price } p^* = p_l. \end{cases} \tag{11}$$

Moreover, by (4) and (5), we have that seller's and the consumers' expected payoffs in equilibrium depend only on the equilibrium price and purchase cascade probabilities, so that in each of the three equilibria we have $\pi^*(p) = \pi^0(p)$ and $C^*(p) = C^0(p)$. It follows that the selection of the undefeated LMSE is also identical to the benchmark case with respect to the equilibrium price. □

Proof of Proposition 6

To prove the result, we show that for the specified range of b there exists $\phi \in (a^2, a)$ such that in the undefeated LMSE, both seller profit and consumer surplus are strictly higher than the benchmarks π_i^0 and C^0. Recall that the benchmark payoffs occur for any $\phi \in [a, 1)$.

When $\phi \to a^-$, there exist four undefeated equilibria, namely, equilibria 1, 3, 5 and 7 in Table 1 (see also Proposition 3). We first show that if $b \in [\beta_l, \beta_m]$, then the undefeated LMSE (i.e., the undefeated equilibrium preferred by the good seller) is equilibrium 5. Equilibria 1 and 3 are identical to the benchmark case. For equilibrium 5, note that when $\phi \to a^-$, we have $p_5 \to \frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)} = p_h$, $\mu \to 0$, $\Gamma_G \to \frac{a^2}{1-a+a^2}$, and therefore the good seller's profit is $\pi_{5G} \to \left(\frac{a^2}{1-a+a^2}\right)\left(\frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)}\right)$. For equilibrium 7, note that when $\phi \to a^-$, we have $p_7 \to \frac{1}{1+\left(\frac{1-a}{a}\right)^2\left(\frac{1-b}{b}\right)}$, $\mu \to 0$, $\Gamma_G \to \frac{a^3}{1-a+a^2}$, and therefore $\pi_{7G} \to \left(\frac{a^3}{1-a+a^2}\right)\left(\frac{1}{1+\left(\frac{1-a}{a}\right)^2\left(\frac{1-b}{b}\right)}\right)$. It follows that $\pi_{7G} \leq \pi_{5G}$ if and only if $b \geq \beta_l = \frac{a^2+a-1}{2a^2+a-1}$, $\pi_{5G} \geq \pi_{3G}$ if and only if $b \leq \beta_m$, and $\pi_{3G} \geq \pi_{1G}$ if and only if $b \leq \beta_h$. Therefore, if $b \in [\beta_l, \beta_m]$ then equilibrium 5 is the undefeated LMSE.

We now return to the proof of the main result. We show that if $b \in [\beta_l, \beta_m]$, there exists $\phi < a$ such that $W^* > W^0$, $\pi_G^* > \pi_G^0$, $\pi_B^* > \pi_B^0$, and $C^* > C^0$. For each of the inequalities, we first establish continuity of the payoff in ϕ at $\phi = a$, and then show that the derivative of the payoff in the limit $\phi \to a^-$ is negative.

Total Welfare. Note that total welfare is given by $W^* = b\Gamma_G$. We have that when $b \in [\beta_l, \beta_m]$ and $\phi \to a^-$, equilibrium 5 is the undefeated LMSE and the good seller's purchase cascade probability $\Gamma_G^* \to \frac{a^2}{1-a+a^2} = \Gamma_G^0$. This implies that the total welfare W^* is continuous at $\phi = a$.

Next, under equilibrium 5, we have $\mu = \frac{a-1+\sqrt{(1-a)^2+\frac{4(a-\phi)}{1-a}}}{2a}$ for $a^2 < \phi < a$. It follows that $\frac{d\mu}{d\phi} = -\frac{1}{2\mu a^2(1-a)+a(1-a)^2} < 0$ and $\lim_{\phi\to a^-} \mu = 0$. The purchase cascade probability for the good seller is given by

$$\begin{aligned}\Gamma_G &= a\Gamma_G(p, b^p) + (1-a)\Gamma_G(p, b^n) \\ &= a\left(\frac{a}{1-a+a^2}\right) + (1-a)\left(\frac{a^2[\mu^2 + a\mu(1-\mu)]}{1-a+a^2}\right),\end{aligned}$$

where the last equality follows from the purchase cascade probabilities derived in Lemma 1 and the fact that under equilibrium 5 we have $b^p \in (b_m, b_h)$ and $b^n = b_l$ (see proof of Proposition 3). Therefore,

$$\begin{aligned}\lim_{\phi\to a^-} \frac{d\Gamma_G}{d\phi} &= \left(\frac{(1-a)a^3}{1-a+a^2}\right) \lim_{\phi\to a^-} \frac{d\mu}{d\phi} \\ &= -\left(\frac{(1-a)a^3}{1-a+a^2}\right) \lim_{\phi\to a^-} \left(\frac{1}{2\mu a^2(1-a)+a(1-a)^2}\right) \\ &= -\left(\frac{(1-a)a^3}{1-a+a^2}\right)\left(\frac{1}{a(1-a)^2}\right) = -\frac{a^2}{(1-a)(1-a+a^2)} < 0.\end{aligned}$$

Therefore, W^* is strictly higher than W^0 for some $\phi \in (a^2, a)$.

Good Seller Profit. Now consider the good seller's profit. Note that under equilibrium 5, we have $p_5 = \frac{1}{1+\left(\frac{1-\phi}{2a-\phi}\right)\left(\frac{1-b}{b}\right)}$ so that $\lim_{\phi\to a^-} p_5 = \frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)} = p_6$. Moreover, from above we have that $\Gamma_G^* \to \Gamma_G^0$, and from Proposition 3 and Lemma 1 we have that for $\phi \geq a$ the undefeated LMSE when $b \in [\beta_l, \beta_m]$ is equilibrium 6; this implies that the good seller's profit function $\pi_G^* = p^*\Gamma_G$ is continuous at $\phi = a$.

Next, consider the derivative of the good seller's profit. Note that $\frac{dp_5}{d\phi} = \frac{b(1-b)(2a-1)}{(b+\phi-2ab-1)^2}$ so that $\lim_{\phi\to a^-} \frac{dp_5}{d\phi} = \frac{b(1-b)(2a-1)}{[ab+(1-a)(1-b)]^2}$. We thus have

$$\begin{aligned}\lim_{\phi\to a^-} \frac{d\pi_G}{d\phi} &= \lim_{\phi\to a^-}\left(\frac{d\Gamma_G}{d\phi}p + \Gamma_G\frac{dp}{d\phi}\right)\\ &= -\frac{a^2 b\left[a^2 b + (1-a)^2(1-b)\right]}{(1-a)\left(1-a+a^2\right)\left[ab+(1-a)(1-b)\right]^2} < 0.\end{aligned}$$

Therefore, π_G^* is strictly higher than π_G^0 for some $\phi \in (a^2, a)$.

Bad Seller Profit. Consider next the profit of a bad seller π_B^*. For a bad seller, we have

$$\lim_{\phi\to a^-} \Gamma_B = \frac{(1-a)^2}{1-a+a^2} = \Gamma_B^0,$$

which implies that in the LMSE, the bad seller's profit function $\pi_B^* = p^*\Gamma_B$ is continuous at $\phi = a$. Furthermore, similar to the good seller profit analysis, it can be shown that under equilibrium 5,

$$\begin{aligned}\lim_{\phi\to a^-} \frac{d\pi_B}{d\phi} &= \lim_{\phi\to a^-}\left(\frac{d\Gamma_B}{d\phi}p + \Gamma_B\frac{dp}{d\phi}\right)\\ &= -\frac{(1-a)b\left[a^2 b + (1-a)^2(1-b)\right]}{\left(1-a+a^2\right)\left[ab+(1-a)(1-b)\right]^2} < 0.\end{aligned}$$

Therefore, π_B^* is strictly higher than π_B^0 for some $\phi \in (a^2, a)$.

Consumer Surplus. Finally, consider the consumers' surplus C^*. Observe that total welfare is given by $W^* = \pi^* + C^*$, and that W^* and π^* have been shown above to be continuous at $\phi = a$, which means that C^* is also continuous at $\phi = a$.

Moreover, we have

$$\begin{aligned}\lim_{\phi\to a^-} \frac{dC}{d\phi} &= \lim_{\phi\to a^-} \frac{dW}{d\phi} - \left(b\lim_{\phi\to a^-}\frac{d\pi_G}{d\phi} + (1-b)\lim_{\phi\to a^-}\frac{d\pi_B}{d\phi}\right)\\ &= -\frac{(2a-1)(1-a)(1-b)[2ab+(1-a)(1-b)]}{(1-a)\left(1-a+a^2\right)\left[ab+(1-a)(1-b)\right]^2} < 0.\end{aligned}$$

Therefore, C^* is strictly higher than C^0 for some $\phi \in (a^2, a)$. □

Proof of Proposition 7

The following lemma is used here as well as in the proof of Proposition 10.

LEMMA 2. *In any PBE of the manipulation game, the seller's value function $\pi_{i,t}^*(b_t)$ is a step function of the belief b_t.*

Proof. The statement holds for any $t > T$ (see the benchmark analysis in §4.1). We prove that it holds also for any $t \leq T$. Note first that in period T we have

$$\pi_{i,T}(b) = P_{i,T}^p \pi_{i,T+1}(b^p) + P_{i,T}^n \pi_{i,T+1}(b^n),$$

where $P^p_{i,T}$ and $P^n_{i,T}$ are the probabilities of a purchase and non-purchase observation occurring in period T, respectively. Note that for any given period-T manipulation strategy (including the equilibrium strategy), these probabilities are independent of b. Since $\pi_{i,T+1}(\cdot)$ is a step function, it follows that $\pi_{i,T}(\cdot)$ is a sum of step functions, and is therefore itself a step function. Applying the same argument reveals that if $\pi_{i,t+1}(\cdot)$ is a step function, then $\pi_{i,t}(\cdot)$ is also a step function. $\square$.

We now return to the proof of the main result. We first show that an equilibrium, if it exists, must be such that in every period $t \le T$ the seller's strategy satisfies $q^*_{G,t} = 0$ and $q^*_{B,t} \in [0,1)$. We then demonstrate that such an equilibrium exists.

We note first that in any period $t \le T$, $q^*_{G,t} = q^*_{B,t} = 1$ cannot be an equilibrium, because this leads to no belief update but a positive probability of getting caught, so that both sellers have an incentive to deviate by decreasing manipulation. We note second that an equilibrium with $q^*_{G,t} \in [0,1)$ and $q^*_{B,t} = 1$ cannot exist, because this leads to $b^n = 1$, so that the bad seller has an incentive to deviate by decreasing manipulation. We note third that an equilibrium with $q^*_{G,t} = 1$ and $q^*_{B,t} \in [0,1)$ cannot exist, because this leads to $b^n = 0$, so that the bad seller has an incentive to deviate by increasing manipulation.

We next show that in any period $t \le T$, an equilibrium with $q^*_{G,t} \in (0,1)$ and $q^*_{B,t} \in (0,1)$ cannot exist. Consider an equilibrium where a real consumer entering the system with belief b purchases if and only if she receives a positive private signal (the analysis of the case where the consumer is indifferent after receiving her private signal follows a similar logic). The updated belief following the observed transaction is

$$\begin{aligned} b^p &= \frac{1}{1 + \left(\frac{1-a+(a-\phi)q_B}{a+(1-a-\phi)q_G}\right)\left(\frac{1-b}{b}\right)}, \\ b^n &= \frac{1}{1 + \left(\frac{1-q_B}{1-q_G}\right)\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}, \end{aligned} \tag{12}$$

for a purchase and non-purchase transaction, respectively. Observe that if $b^p < b^n$, then $\frac{d\pi_{i,t}}{dq_{i,t}} < 0$ for $i \in \{G, B\}$ so that $q_{G,t} = q_{B,t} = 0$ is the only possible equilibrium. However, if $q_{G,t} = q_{B,t} = 0$, then according to (12) we have $b^p > b^n$, resulting in a contradiction. Therefore, an equilibrium, if it exists, must be such that $b^p \ge b^n$. Moreover, note that if $\phi \ge 1-a$, then $\frac{d\pi_{G,t}}{dq_{G,t}} < 0$ so that only $q^*_{G,t} = 0$ is possible in equilibrium. Therefore, in what follows it will suffice to restrict attention to cases of $\phi < 1-a$ and to equilibria satisfying $b^p \ge b^n$.

Observe from (12) that for a given q_B, b^p is increasing in q_G, while b^n is decreasing in q_G. Therefore, for any given belief b, we have that $\frac{d\pi_{G,t}}{dq_{G,t}}$ is increasing in $q_{G,t}$, and an equilibrium exists if and only if $\left.\frac{d\pi_{G,t}}{dq_{G,t}}\right|_{q_{G,t}=0} \le 0$. Assuming this holds, there are two possible equilibria: (1) $\left.\frac{d\pi_{G,t}(b)}{dq_{G,t}}\right|_{q_{G,t}=0} \le 0$ so that $q^*_{G,t} = 0$, or (2) $\left.\frac{d\pi_{G,t}(b)}{dq_{G,t}}\right|_{q_{G,t}=q^*_{G,t}} = 0$ with $q^*_{G,t} > 0$. Note that since $\frac{d\pi_{G,t}}{dq_{G,t}}$ is a step function in $q_{G,t}$, the latter case in general requires specifying the consumers' randomization strategy $\mu(\cdot)$ appropriately, as in our main analysis (see proof of Proposition 3).

Next, observe also from (12) that for a given q_G, b^p is decreasing in q_B, while b^n is increasing in q_B. Moreover, in the limit $q_{B,t} = 1$, we have $b^n = 1$. Therefore, for any given belief b, $\frac{d\pi_{B,t}}{dq_{B,t}}$ is decreasing in $q_{B,t}$. We note also that $\left.\frac{d\pi_{B,t}(b)}{dq_{B,t}}\right|_{q_{B,t}=1} < 0$, so that if $\left.\frac{d\pi_{B,t}}{dq_{B,t}}\right|_{q_{B,t}=0} < 0$ then $\frac{d\pi_{B,t}}{dq_{B,t}} < 0$ for any $q_{B,t} \in [0,1]$. In this case, the only possible equilibrium is $q^*_{B,t} = 0$. On the other hand, if $\left.\frac{d\pi_{B,t}}{dq_{B,t}}\right|_{q_{B,t}=0} > 0$, then an equilibrium, if it exists, must be such that $q^*_{B,t} > 0$. Since $\frac{d\pi_{B,t}}{dq_{B,t}}$ is a step function in $q_{B,t}$, the latter case requires specifying the consumers' randomization strategy $\mu(\cdot)$ appropriately.

Combining the above two paragraphs, observe that an equilibrium with $q^*_{G,t} > 0$ and $q^*_{B,t} > 0$ in general cannot exist. To see this, notice that $\frac{d\pi_{G,t}}{dq_{G,t}}$ and $\frac{d\pi_{B,t}}{dq_{B,t}}$ are step functions with different step levels, so that specifying a consumer randomization strategy $\mu(\cdot)$ which simultaneously satisfies both $\frac{d\pi_{G,t}}{dq_{G,t}} = 0$ and $\frac{d\pi_{B,t}}{dq_{B,t}} = 0$ is impossible. In particular, we note that indifference of both sellers at positive manipulation values requires

$$(1-a-\phi)\pi_{G,t+1}(b^p) - (1-a)\pi_{G,t+1}(b^n) = 0, \text{ and}$$
$$(a-\phi)\pi_{B,t+1}(b^p) - a\pi_{B,t+1}(b^n) = 0.$$

While the consumers' randomization strategy $\mu(\cdot)$ can be specified such that one of the above two equations is satisfied, it is in general impossible to specify $\mu(\cdot)$ such that both are satisfied (note that $\mu(\cdot)$ does not depend on the seller's type). It follows that an equilibrium with $q^*_{G,t} > 0$ and $q^*_{B,t} > 0$ cannot exist.

We now show that an equilibrium with $q^*_{G,t} > 0$ and $q^*_{B,t} = 0$ also cannot exist. Suppose that such an equilibrium exists, which implies $\frac{d\pi_{G,t}}{dq_{G,t}} = 0$ and $\frac{d\pi_{B,t}}{dq_{B,t}} \leq 0$. Consider first the case $b < p$. If $q^*_{G,t} > 0$ and $q^*_{B,t} = 0$, then using (12) we have $b^n < b_l$, which implies $\pi_{G,t+1}(b^n) = \pi_{B,t+1}(b^n) = 0$. However, this leads to $\frac{d\pi_{G,t}}{dq_{G,t}} > 0$ and $\frac{d\pi_{B,t}}{dq_{B,t}} > 0$, resulting in a contradiction. Now consider the case $b \geq p$. If $q^*_{G,t} > 0$ and $q^*_{B,t} = 0$, using (12) we have $b^p > b_h$ and $\pi_{G,t+1}(b^p) = \pi_{B,t+1}(b^p) = p$. However, this leads to

$$\begin{aligned}
\frac{d\pi_{B,t}}{dq_{B,t}} &= (a-\phi)\pi_{B,t+1}(b^p) - a\pi_{B,t+1}(b^n) \\
&= (a-\phi)p - a\pi_{B,t+1}(b^n) \\
&> (1-a-\phi)p - (1-a)\pi_{B,t+1}(b^n) \\
&\geq (1-a-\phi)p - (1-a)\pi_{G,t+1}(b^n) \\
&= \frac{d\pi_{G,t}}{dq_{G,t}} = 0,
\end{aligned}$$

which contradicts the equilibrium condition $\frac{d\pi_{B,t}}{dq_{B,t}} \leq 0$. Therefore, an equilibrium with $q^*_{G,t} > 0$ and $q^*_{B,t} = 0$ cannot exist.

From the above discussion, it follows that an equilibrium, if it exists, must be such that in any period $t \leq T$, we have $q^*_{G,t} = 0$ and $q^*_{B,t} \in [0,1)$. Note that Proposition 10 establishes the existence and properties of $q^*_{B,t} \in [0,1)$, assuming that the good seller never manipulates (i.e., assuming $q^*_{G,t} = 0$). Here, we show that when the bad seller's strategy is $q^*_{B,t} \in [0,1)$ as described in Proposition 10, $q^*_{G,t} = 0$ is indeed an equilibrium strategy for the good seller. We prove the result by backward induction. The claim holds in period T (the period-T equilibrium is identical to the manipulation equilibrium in our main model); suppose the claim holds in period $t+1$, and consider period t.

Note first that in period t, we have for the bad seller the equilibrium condition

$$\frac{d\pi_{B,t}}{dq_{B,t}} = (a-\phi)\pi_{B,t+1}(b^p) - a\pi_{B,t+1}(b^n) \leq 0,$$

and recall that $q^*_{G,t+1} = 0$ by the inductive hypothesis. Then, there are three (exhaustive) cases to consider, based on the bad seller's equilibrium strategy in period $t+1$, $q^*_{B,t+1}$.

Case (i): The equilibrium is such that $q^*_{B,t+1}(b^n)=0$ and $q^*_{B,t+1}(b^p)\geq 0$. Suppose first that in equilibrium we have $b^p>b_h$. Then, we have for the good seller

$$\begin{aligned}\frac{d\pi_{G,t}}{dq_{G,t}} &= (1-a-\phi)\pi_{G,t+1}(b^p)-(1-a)\pi_{G,t+1}(b^n)\\ &= (1-a-\phi)p-(1-a)\pi_{G,t+1}(b^n)\\ &< (a-\phi)p-a\pi_{G,t+1}(b^n)\\ &\leq (a-\phi)p-a\pi_{B,t+1}(b^n)\\ &= \frac{d\pi_{B,t}}{dq_{B,t}}\leq 0,\end{aligned}$$

so that an equilibrium with $q^*_{G,t}=0$ exists. Next suppose that in equilibrium we have $b^p\leq b_h$, and note that this implies that $b^{pn}<b_h$. If $q^*_{G,t+1}(b^n)=q^*_{B,t+1}(b^n)=0$, it must be that $b^n\geq p$ (otherwise, the bad seller has an incentive to deviate by increasing manipulation). Moreover, if $q^*_{G,t+1}(b^n)=q^*_{B,t+1}(b^n)=0$ and $b^n\geq p$, then it follows by (12) that $b^{np}\geq b_h$. Therefore, we have $b^{np}>b^{pn}$, which implies $\pi_{G,t+2}(b^{np})\geq\pi_{G,t+2}(b^{pn})$. The good seller's derivative is

$$\begin{aligned}\frac{d\pi_{G,t}}{dq_{G,t}} &= (1-a-\phi)\pi_{G,t+1}(b^p)-(1-a)\pi_{G,t+1}(b^n)\\ &= (1-a-\phi)\left[a\pi_{G,t+2}(b^{pp})+(1-a)\pi_{G,t+2}(b^{pn})\right]-(1-a)\left[a\pi_{G,t+2}(b^{np})+(1-a)\pi_{G,t+2}(b^{nn})\right]\\ &= a\left[(1-a-\phi)\pi_{G,t+2}(b^{pp})-(1-a)\pi_{G,t+2}(b^{pn})\right]+(1-\phi)(1-a)\pi_{G,t+2}(b^{pn})\\ &\quad (1-a)\left[(1-a-\phi)\pi_{G,t+2}(b^{np})-(1-a)\pi_{G,t+2}(b^{nn})\right]-(1-\phi)(1-a)\pi_{G,t+2}(b^{np})\\ &= a\left.\frac{d\pi_{G,t+1}(b^p)}{dq_{G,t+1}=0}\right|_{q_{G,t+1}=0}+(1-a)\left.\frac{d\pi_{G,t+1}(b^n)}{dq_{G,t+1}}\right|_{q_{G,t+1}=0}\\ &\quad +(1-\phi)(1-a)\left[\pi_{G,t+2}(b^{pn})-\pi_{G,t+2}(b^{np})\right]\\ &\leq a\left.\frac{d\pi_{G,t+1}(b^p)}{dq_{G,t+1}}\right|_{q_{G,t+1}=0}+(1-a)\left.\frac{d\pi_{G,t+1}(b^n)}{dq_{G,t+1}}\right|_{q_{G,t+1}=0}\leq 0,\end{aligned}$$

where the second-to-last inequality follows from $\pi_{G,t+2}(b^{np})\geq\pi_{G,t+2}(b^{pn})$ and the last inequality follows from the inductive hypothesis (which implies $\left.\frac{d\pi_{G,t+1}}{dq_{G,t+1}}\right|_{q_{G,t+1}=0}\leq 0$). Therefore, an equilibrium with $q^*_{G,t}=0$ exists.

Case (ii): The equilibrium is such that $q^*_{B,t+1}(b^n)>0$ and $q^*_{B,t+1}(b^p)=0$. Note that if $q_{G,t+1}(b^p)=q_{B,t+1}(b^p)=0$, it must be that $b^p\geq p$ (otherwise the bad seller would have an incentive to deviate). It then follows from (12) that $b_l\leq b^{pn}<p$. Moreover, observe that in this case, it must be that $\pi_{B,t+2}(b^{np})\geq\pi_{B,t+2}(b^{pn})$. [*Proof: Note first that if* $t=T-1$, *then* $b^{np}\geq p$ *because there is no manipulation after period* T. *Therefore, we have* $b^{np}>b^{pn}$, *which implies* $\pi_{B,t+2}(b^{np})\geq\pi_{B,t+2}(b^{pn})$. *Next consider any* $t<T-1$. *Suppose that the claim does not hold. Then it must be that* $b^{np}<b^{pn}<p$, *which also implies* $b^n<p$. *Moreover, note that at each of the posterior beliefs* b^n, b^{np} *and* b^{nn}, *there must be positive manipulation in equilibrium (otherwise, the bad seller's derivative would be positive, contradicting the equilibrium); this implies that* $\pi_{B,t+2}(b^{np})=(1-\phi)\pi_{B,t+3}(b^{npp})$ *and* $\pi_{B,t+2}(b^{nn})=(1-\phi)\pi_{B,t+3}(b^{nnp})$. *Then, from the bad seller's indifference condition at posterior* b^n *and we have*

$$(a-\phi)\pi_{B,t+2}(b^{np})-a\pi_{B,t+2}(b^{nn})=0\iff(a-\phi)(1-\phi)\pi_{B,t+3}(b^{npp})-a(1-\phi)\pi_{B,t+3}(b^{nnp})=0,$$

or, equivalently, $(a-\phi)\pi_{B,t+3}(b^{npp}) - a\pi_{B,t+3}(b^{nnp}) = 0$*. Therefore, in such an equilibrium, the following three equations must be satisfied simultaneously:*

$$(a-\phi)\pi_{B,t+3}(b^{npp}) - a\pi_{B,t+3}(b^{nnp}) = 0,$$
$$(a-\phi)\pi_{B,t+3}(b^{npp}) - a\pi_{B,t+3}(b^{npn}) = 0,$$
$$(a-\phi)\pi_{B,t+3}(b^{nnp}) - a\pi_{B,t+3}(b^{nnn}) = 0,$$

where the last two equations follow from the bad seller's indifference condition when the belief at the beginning of period $t+2$ *is* b^{np} *and* b^{nn}*, respectively. However, observe that there are only two consumer randomizations strategies* $\mu(b^{np})$ *and* $\mu(b^{nn})$ *that enter the above equations, so that the equations in general cannot be satisfied simultaneously.*] Following the same argument as in case (i) above, it can then be shown that $\left.\frac{d\pi_{G,t}}{dq_{G,t}}\right|_{q_{G,t}=0} \leq 0$, which implies that there exists an equilibrium with $q^*_{G,t} = 0$.

Case (iii): The equilibrium is such that $q^*_{B,t+1}(b^n) > 0$ and $q^*_{B,t+1}(b^p) > 0$. Note that this implies $\pi_{B,t+1}(b^p) = (1-\phi)\pi_{B,t+2}(b^{pp})$ and $\pi_{B,t+1}(b^n) = (1-\phi)\pi_{B,t+2}(b^{np})$. Substituting these equations into the bad seller's period-t indifference condition, we have

$$(a-\phi)(1-\phi)\pi_{B,t+2}(b^{pp}) - a(1-\phi)\pi_{B,t+2}(b^{np}) \leq 0 \iff (a-\phi)\pi_{B,t+2}(b^{pp}) - a\pi_{B,t+2}(b^{np}) \leq 0.$$

Moreover, the period-$t+1$ indifference condition at belief b^p is

$$(a-\phi)\pi_{B,t+2}(b^{pp}) - a\pi_{B,t+2}(b^{pn}) = 0.$$

It follows that in such an equilibrium we have $\pi_{B,t+2}(b^{np}) \geq \pi_{B,t+2}(b^{pn})$, which implies also that $\pi_{G,t+2}(b^{np}) \geq \pi_{G,t+2}(b^{pn})$ (since the value functions of both sellers are non-decreasing step functions with discontinuities at the same belief values). The good seller's derivative evaluated at $q_{G,t} = 0$ is

$$\begin{aligned}
\left.\frac{d\pi_{G,t}(b)}{dq_{G,t}}\right|_{q_{G,t}=0} &= (1-a-\phi)\pi_{G,t+1}(b^p) - (1-a)\pi_{G,t+1}(b^n) \\
&= (1-a-\phi)\left[a\pi_{G,t+2}(b^{pp}) + (1-a)\pi_{G,t+2}(b^{pn})\right] - (1-a)\left[a\pi_{G,t+2}(b^{np}) + (1-a)\pi_{G,t+2}(b^{nn})\right] \\
&= a\left[(1-a-\phi)\pi_{G,t+2}(b^{pp}) - (1-a)\pi_{G,t+2}(b^{pn})\right] + (1-\phi)(1-a)\pi_{G,t+2}(b^{pn}) \\
&\quad (1-a)\left[(1-a-\phi)\pi_{G,t+2}(b^{np}) - (1-a)\pi_{G,t+2}(b^{nn})\right] - (1-\phi)(1-a)\pi_{G,t+2}(b^{np}) \\
&\leq a\left.\frac{d\pi_{G,t+1}(b^p)}{dq_{G,t+1} = 0}\right|_{q_{G,t+1}=0} + (1-a)\left.\frac{d\pi_{G,t+1}(b^n)}{dq_{G,t+1}}\right|_{q_{G,t+1}=0} \leq 0,
\end{aligned}$$

where the last inequality follows from the inductive hypothesis (which implies $\left.\frac{d\pi_{G,t+1}(\cdot)}{dq_{G,t+1}}\right|_{q_{G,t+1}=0} \leq 0$). Therefore, there exists an equilibrium with $q^*_{G,t} = 0$. This completes the induction proof.

Proof of Proposition 8

Fix any price p. The value function $\pi_{i,t}(b)$ in period $t = T+1$ is equal to the benchmark case without manipulation. In period $t = T$, the manipulation game is equivalent to the manipulation game in our main model (i.e., our main model corresponds to $T = 1$). By Proposition 2 and Corollary 1, if $\phi \geq a$ then $q^*_{i,T}(b) = 0$ for all b. It follows that $\pi_{i,T}(b) = \pi_{i,T+1}(b) = \pi^0_i(b)$. Repeating the argument establishes the result. □

Proof of Proposition 9

Fix any price p. The value function $\pi_{i,t}(b)$ in period $t = T + 1$ is equal to the benchmark case without manipulation. In period $t = T$, the manipulation game is equivalent to the manipulation game in our main model (i.e., our main model corresponds to $T = 1$). By the proof of Proposition 5, it follows that in any undefeated equilibrium, we have $\pi^*_{i,T}(b) \to \pi^*_{i,T+1}(b) = \pi^0_i(b)$ and $C^*_T(b) \to C^*_{T+1}(b) = C^0(b)$. Repeating the argument establishes the result. □

Proof of Proposition 10

By Proposition 7, we fix $q_{G,t} = 0$ and focus on establishing the existence and properties of the bad seller's period-t equilibrium manipulation strategy. For ease of exposition, we suppress the subscript "B" in the bad seller's profit function and manipulation strategy. By Lemma 2, the bad seller's derivative is given by

$$\frac{d\pi_t}{dq_t} = (a - \phi)\pi_{t+1}(b^p) - a\pi_{t+1}(b^n)$$

Recall that b^p (b^n) is strictly decreasing (strictly increasing) in the bad seller's manipulation strategy q_t. Since $\pi_{t+1}(b)$ is a non-decreasing step function of b, the bad seller's derivative is a non-increasing step function of q_t. Moreover, we note that when $q_t = 1$ we have $b^n > b^p$, which implies that the above derivative is negative at $q_t = 1$. Then, there are two possibilities: (i) the derivative at $q_t = 0$ is nonpositive and therefore an equilibrium with $q^*_t = 0$ exists, or (ii) the derivative at $q_t = 0$ is positive and an equilibrium at $q^*_t > 0$ exists. As in our main analysis, equilibrium in the latter case in general requires specifying the consumers' randomization strategy $\mu(\cdot)$ appropriately (see proof of Proposition 3).

We next prove the statements made in the proposition. We first prove part (i) by contradiction. Suppose $b_t = b < p$ and the bad seller's equilibrium strategy is $q^*_t = 0$. If a non-purchase observation occurs in period t, the posterior belief is $b^n = \frac{1}{1+\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)} < \frac{1}{1+\left(\frac{a}{1-a}\right)\left(\frac{1-p}{p}\right)}$, and a non-purchase cascade is triggered. However, this cannot be an equilibrium because the bad seller's derivative is

$$\left.\frac{d\pi_t(b)}{dq_t}\right|_{q_t=0} = (a - \phi)\pi_{t+1}(b^p) - a\pi_{t+1}(b^n) = (a - \phi)\pi_{t+1}(b^p) > 0,$$

which implies that the bad seller has a profitable deviation.

We next prove part (ii) of the proposition. Suppose first that $b_t = b > p$ and $q^*_t = 0$. If a purchase observation occurs in period t, the posterior belief is $b^p = \frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-b}{b}\right)} > \frac{1}{1+\left(\frac{1-a}{a}\right)\left(\frac{1-p}{p}\right)}$, which implies that a purchase cascade is triggered and $\pi_{t+1}(b^p) = p$. The bad seller's derivative is

$$\left.\frac{d\pi_t(b)}{dq_t}\right|_{q_t=0} = (a - \phi)\pi_{t+1}(b^p) - a\pi_{t+1}(b^n) = (a - \phi)p - a\pi_{t+1}(b^n),$$

where $b^n = \frac{1}{1+\left(\frac{a}{1-a}\right)\left(\frac{1-b}{b}\right)}$. Since $b > p$, we have $b^n < p$ and from statement (i) above, we know that $q^*_{t+1} > 0$. At the same time, indifference of the bad seller in period $t + 1$ implies that $(a - \phi)\pi_{t+2}(b^{np}) = a\pi_{t+2}(b^{nn})$, which further implies that

$$\pi_{t+1}(b^n) = (1 - \phi)\pi_{t+2}(b^{np}).$$

Now, observe that if $q^*_{t+2}(b^{np}) > 0$, we may further write

$$\pi_{t+1}(b^n) = (1 - \phi)^2\pi_{t+3}(b^{npp}),$$

while if $q^*_{t+2}(b^{np}) = 0$, we may write

$$\pi_{t+1}(b^n) = (1-\phi)\left[(1-a)\pi_{t+3}(b^{npp}) + a\pi_{t+3}(b^{npn})\right].$$

Repeating the above process up to period $T+1$ and using the fact that $\pi_{T+1}(\cdot)$ is piecewise constant and independent of ϕ, we conclude that $\pi_{t+1}(b^n)$ can be expressed as

$$\pi_{t+1}(b^n) = \sum_{i=1}^{T-t} c_i(1-\phi)^i,$$

where $c_i \geq 0$ for $1 \leq i \leq T-t$. Next, let

$$f_t(\phi) = (a-\phi)p - a\pi_{t+1}(b^n) = (a-\phi)p - a\sum_{i=1}^{T-t} c_i(1-\phi)^i.$$

Note that $f_t''(\phi) \leq 0$, that is, $f_t(\phi)$ is concave in ϕ. Note also that $f_t(0) > 0$ because $p > \pi_{t+1}(b^n)$, and $f_t(a) < 0$ because $\pi_{t+1}(b^n) > 0$. It follows that $f_t(\phi) = 0$ admits a unique root in the interval $\phi \in (0,a)$. Let $\rho_t(b)$ be the solution to $f_t(\phi) = 0$. We then have $\frac{d\pi_t(b)}{dq_t}\Big|_{q_t=0} = f_t(\phi) > 0$ if and only if $\phi < \rho_t(b)$, which implies that $q_t^* > 0$ if and only if $\phi < \rho_t(b)$.

Now suppose $b = p$ and $q_t^* = 0$. If a purchase (non-purchase) observation occurs in period t, the posterior belief is $b^p = b_h$ ($b^n = b_l$). The bad seller's derivative at $q_t = 0$ is

$$\frac{d\pi_t(b)}{dq_t}\bigg|_{q_t=0} = (a-\phi)\pi_{t+1}(b^p) - a\pi_{t+1}(b^n) = (a-\phi)\pi_{t+1}(b_h) - a\pi_{t+1}(b_l).$$

Note that $b_{t+1} = b_h > p$, and the result of part (ii) applies. That is, there exists $\rho_{t+1}(b_h)$ such that $q^*_{t+1} > 0$ if and only if $\phi < \rho_{t+1}$.

Consider first the case of $\phi \geq \rho_{t+1}$, where we have $q^*_{t+1} = 0$ for $b_{t+1} = b_h$. The bad seller's profit function is

$$\pi_{t+1}(b_h) = (1 - a + a\mu_{t+1})p + a(1-\mu_{t+1})\pi_{t+2}(p) \leq p,$$

where equality holds for $\mu_{t+1} = 1$. For the case of $\mu_{t+1} = 1$, using the same argument as in part (ii) above, we consider

$$g(\phi) = (a-\phi)p - a\pi_{t+1}(b_l) = (a-\phi)p - a\sum_{i=1}^{T-t} \tilde{c}_i(1-\phi)^i.$$

It can be shown that $g(\phi) = 0$ admits a unique root in the interval $\phi \in (0,a)$; let this root be ϕ_0. For any $\phi \in [\phi_0, a)$, we have

$$\frac{d\pi_t(b)}{dq_t}\bigg|_{q_t=0} = (a-\phi)\pi_{t+1}(b_h) - a\pi_{t+1}(b_l) \leq (a-\phi)p - a\pi_{t+1}(b_l) = g(\phi) \leq 0.$$

Therefore, there exists a threshold $\overline{\rho}_t(b) \leq \phi_0$ such that $q_t^* = 0$ for $\phi \geq \overline{\rho}_t(b)$.

Consider next the case of $\phi < \rho_{t+1}(b_h)$. We have $q^*_{t+1} > 0$ for $b_{t+1} = b_h$, which implies $\pi_{t+1}(b_h) = (1-\phi)p$. Consider

$$g(\phi) = (a-\phi)(1-\phi)p - a\pi_{t+1}(b_l) = (1-\phi)\left[(a-\phi)p - a\sum_{i=1}^{T-t} \tilde{c}_i(1-\phi)^{i-1}\right] = (1-\phi)\tilde{g}(\phi).$$

It can be shown that $\tilde{g}(\phi) = 0$ admits a unique root in the interval $\phi \in (0,a)$; let this root be ϕ_1. Moreover, note that when $b_{t+1} = b_h$ and $\phi < \rho_{t+1}(b_h)$,

$$\frac{d\pi_{t+1}(b_{t+1})}{dq_{t+1}}\bigg|_{q_{t+1}=0} = (a-\phi)p - a\pi_{t+2}(p) > 0.$$

Observe that $b_l < p$. By part (i), we have $q_{t+1} > 0$ in this case. Thus, $\pi_{t+1}(b_l) = (1-\phi)\pi_{t+2}(b_l^p) \leq (1-\phi)\pi_{t+2}(p)$, where we use the fact that $b_l^p \leq p$. It follows that if $\phi < \rho_{t+1}(b_h)$, we have

$$\begin{aligned} g(\phi) &= (a-\phi)(1-\phi)p - a\pi_{t+1}(b_l) \\ &\geq (a-\phi)(1-\phi)p - a(1-\phi)\pi_{t+2}(p) \\ &= (1-\phi)\left[(a-\phi)p - a\pi_{t+2}(p)\right] > 0, \end{aligned}$$

which implies that $\rho_{t+1}(b_h) < \phi_1$. Therefore, there exists a threshold $\underline{\rho}_t(b) \geq \rho_{t+1}(b_h)$ such that $q_t^* > 0$ for $\phi < \underline{\rho}_t(b)$. □

Acknowledgements

The authors thank the Department Editor, the Associate Editor, and three anonymous referees for their constructive comments and suggestions. The authors are also grateful to Justin Johnson, Michael Waldman, and seminar participants at Boston College, the University of Southern California, the University of Chicago, the Utah Winter Operations Conference, INFORMS 2019 and M&SOM 2019 conferences for their helpful feedback.

References

Acemoglu, D., M. Dahleh, I. Lobel, A. Ozdaglar. 2011. Bayesian learning in social networks. *The Review of Economic Studies* **78**(4) 1201–1236.

Banerjee, A.V. 1992. A simple model of herd behavior. *The Quarterly Journal of Economics* **107**(3) 797–817.

BBC. 2017. Celebrities warned over instagram ads. (Apr. 20th).

Bikhchandani, S., D. Hirshleifer, I. Welch. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* **100**(5) 992–1026.

Bimpikis, K., Y. Papanastasiou, W. Zhang. 2019. Information disclosure in service platforms: Optimizing for supply *Working Paper, Stanfrod University.*

Chakraborty, S., R. Swinney. 2017. Signaling to the crowd: Private quality information and rewards-based crowdfunding. *Working Paper, Duke University.*

Che, Yeon-Koo, Johannes Hörner. 2018. Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics* **133**(2) 871–925.

Chen, Y.-J., T. Dai, G. Korpeoglu, E. Körpeoğlu, O. Sahin, C. Tang, S. Xiao. 2019. Innovative online platforms: Research opportunities. *Manufacturing & Service Operations Management* .

Cho, I.-K., D. Kreps. 1987. Signaling games and stable equilibria. *The Quarterly Journal of Economics* **102**(2) 179–221.

Cho, S.-H., X. Fang, S. Tayur. 2015. Combating strategic counterfeiters in licit and illicit supply chains. *Manufacturing & Service Operations Management* **17**(3) 273–289.

Crapis, D., B. Ifrach, C. Maglaras, M. Scarsini. 2017. Monopoly pricing in the presence of social learning *Management Science, forthcoming.*

Debo, L., C. Parlour, U. Rajan. 2012. Signaling quality via queues. *Management Science* **58**(5) 876–891.

Debo, L., U. Rajan, S. Veeraraghavan. 2019. Signaling quality via long lines and uninformative prices. *Manufacturing & Service Operations Management* .

Drakopoulos, K., A. Ozdaglar, J. Tsitsiklis. 2013. On learning with finite memory. *IEEE Transactions on Information Theory* **59**(10) 6859–6872.

Federal Trade Commission. 2016. Lord & taylor settles ftc charges it deceived consumers through paid article in an online fashion magazine and paid instagram posts by 50 "fashion influencers". (Mar. 15th).

Federal Trade Commission. 2019. FTC brings first case challenging fake paid reviews on an independent retail website. (Feb. 26th).

Feldman, P., Y. Papanastasiou, E. Segev. 2016. social learning and the design of new experience goods *Working Paper, University of California, Berkeley.*

Financial Times. 2016. China's ecommerce sites try to sweep away "brushing". (Nov. 23rd).

Forbes. 2015. JD.com fires employee for encouraging 'brushing' fake orders to inflate sales. (Apr. 11th).

Forbes. 2019. Amazon's fake review problem is getting worse. (Apr. 16th).

Fortune. 2016. Amazon is cracking down on more fake reviews. (Oct. 27th).

Fortune. 2019. Amazon prime day: Beware, fake reviews may run rampant. (Jul. 15th).

Frazier, P., D. Kempe, J. Kleinberg, R. Kleinberg. 2014. Incentivizing exploration. *Technical Report, Cornell University.*

Gao, Y., W. Lim, C. Tang. 2016. Entry of copycats of luxury brands. *Marketing Science* **36**(2) 272–289.

Godes, D. 2016. Product policy in markets with word-of-mouth communication. *Management Science* .

Guardian. 2017. Social media stars face crackdown over money from brandsd. (Sep. 16th).

Guarino, A., H. Harmgart, S. Huck. 2011. Aggregate information cascades. *Games and Economic Behavior* **73**(1) 167–185.

Herrera, H., J. Hörner. 2013. Biased social learning. *Games and Economic Behavior* **80** 131–146.

Hu, M., J. Milner, J. Wu. 2015. Liking and following and the newsvendor: Operations and marketing policies under social influence. *Management Science* **62**(3) 867–879.

Ifrach, B., C. Maglaras, M. Scarsini, A. Zseleva. 2019. Bayesian social learning from consumer reviews. *Operations Research* **67**(5) 1209–1221.

Jin, C., L. Yang, K. Hosanagar. 2019. To brush or not to brush: Product rankings, customer search, and fake orders. *Working Paper* .

Johnson, J., D.D. Sokol. 2020. Understanding AI collusion and compliance. D.D. Sokol, B. van Rooij, eds., *Cambridge Handbook of Compliance*. Cambridge University Press, Cambridge, UK. Forthcoming.

Kremer, M., L. Debo. 2015. Inferring quality from wait time. *Management Science* **62**(10) 3023–3038.

Luca, Michael, Georgios Zervas. 2016. Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science* **62**(12) 3412–3427.

Mailath, G., M. Okuno-Fujiwara, A. Postlewaite. 1993. Belief-based refinements in signalling games. *Journal of Economic Theory* **60**(2) 241–276.

Mayzlin, Dina, Yaniv Dover, Judith Chevalier. 2014. Promotional reviews: An empirical investigation of online review manipulation. *American Economic Review* **104**(8) 2421–55.

Papanastasiou, Y. 2019a. Fake news propagation and detection: A sequential model *Management Science, forthcoming.*

Papanastasiou, Y. 2019b. Newsvendor decisions with two-sided learning *Management Science, forthcoming.*

Papanastasiou, Y., K. Bimpikis, N. Savva. 2017. Crowdsourcing exploration *Management Science, forthcoming.*

Papanastasiou, Y., N. Savva. 2016. Dynamic pricing in the presence of social learningand strategic consumers. *Management Science* **63**(4) 919–939.

Pun, H., G. DeYong. 2017. Competing with copycats when customers are strategic. *Manufacturing & Service Operations Management* **19**(3) 403–418.

Pun, H., J. Swaminathan, P. Hou. 2018. Blockchain adoption for combating deceptive counterfeits. *Working Paper* .

Qian, Y., Q. Gong, Y. Chen. 2014. Untangling searchable and experiential quality responses to counterfeits. *Marketing Science* **34**(4) 522–538.

Schmidt, W., V. Gaur, R. Lai, A. Raman. 2015. Signaling to partially informed investors in the newsvendor model. *Production and Operations Management* **24**(3) 383–401.

Shin, D., A. Zeevi. 2017. Dynamic pricing and learning with online product reviews. *Working paper, Columbia University* .

Smith, L., P. Sørensen. 2000. Pathological outcomes of observational learning. *Econometrica* **68**(2) 371–398.

The Wall Street Journal. 2015. They call it "brushing": The dark art of alibaba sales fakery. (Mar. 3rd).

The Wall Street Journal. 2017. Social-media stars are turning heads—of regulators. (Dec. 15th).

Veeraraghavan, S., L. Debo. 2009. Joining longer queues: Information externalities in queue choice. *Manufacturing & Service Operations Management* **11**(4) 543–562.

Veeraraghavan, S.K., L.G. Debo. 2011. Herding in queues with waiting costs: Rationality and regret. *Manufacturing & Service Operations Management* **13**(3) 329–346.

Yu, M., L. Debo, R. Kapuscinski. 2015. Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods. *Management Science* **62**(2) 410–435.