

Recommendations with Feedback*

Ganesh Iyer

Gustavo Manso

University of California, Berkeley

University of California, Berkeley

This version: February 2019

*We thank Jonathan Berk, Ernesto Dal Bo, Brett Green, Yuichiro Kamada, Abhishek Nagaraj, Steve Tadelis, Boris Vallee and seminar participants at Boston College, Berkeley Haas Faculty seminar, Berkeley-Stanford Finance seminar, and University of Maryland for comments. Address for correspondence: Haas School of Business, University of California, Berkeley, Berkeley, CA 94720-1900. Email: giyer@haas.berkeley.edu, manso@haas.berkeley.edu.

RECOMMENDATIONS WITH FEEDBACK

ABSTRACT

We investigate the strategic role of a recommender who cares about accuracy and whose recommendations influence product quality. In the presence of such feedback effects, recommendations have a self-fulfilling property: the recommendation agent can select any firm which will end up being the firm with the best quality. Recommendations can lead to significant inefficiencies which include: i) a lack of incentive to acquire valuable information, ii) a status quo bias, and iii) the avoidance of risky innovations. Direct monetary payments from firms may work in mitigating these inefficiencies, while competition between recommenders and monetary transfers from consumers are ineffective.

I Introduction

In a broad range of markets there is uncertainty about which firm will succeed in innovation and will end up having the highest quality product. But whether a firm can successfully offer the best product may in turn depend upon whether a significant number of buyers somehow decide to take a chance on its product being successful. If many buyers were to indeed buy the product under the expectation that it will be successful, then the firm will have the resources to develop a successful high quality product. Yet unless buyers are convinced that the product will be successful, they have no reason to adopt the product, suggesting a coordination problem. In this paper we study the strategic role of recommendation agents or opinion leaders in resolving these types of coordination problems.

An important aspect of the success of recommendation agents as coordination devices is the existence of feedback mechanisms. For example, annual business school rankings by *Businessweek*, the *Economist*, and the *Financial Times* affect a school's ability to attract high quality applicants and to charge full tuition. Many deans, however, also recognize some other equally important benefits of rankings. Namely, better rankings may induce status conscious donors to contribute, help to attract better quality faculty, or attract higher paying recruiters.¹ These benefits can act as "feedback" mechanisms: Greater donations or better faculty talent can endogenously increase the actual quality of the school making it consistent with the expectations of applicants (consumers) who were swayed by the rankings in the first place. In this manner recommendations by influential agents can become self-fulfilling prophecies.

The phenomenon described above has substantial economic significance. There are many other important examples in which recommendations help resolve coordination problems between firms and consumers through the presence of a feedback channel:

- In the wine industry Robert Parker is one of the most widely known and influential

¹See Gellman, "Do Business School Rankings Matter?" *The Wall Street Journal*, 03/04/2015. Further, in a study based on the *US News and World Report* law school rankings Espeland and Sauder (2007) find that rankings impact schools through their effects not only on prospective students, but also on other audiences such as trustees, board of visitors, and alumni. In addition, rankings also affect the amount of funding a school can secure from the larger university budget.

critic. His 100 point rating system in the *The Wine Advocate* has become influential in American wine buying and is particularly vital for the success of newly released Bordeaux wines. No merchant or wine investor dared ignore him. His positive call on the 1982 Bordeaux vintage was crucial for its success and his 100 point rating of Chateau Haut-Brion 1989 means that it commands a significant premium on other Haut-Brion vintages. Importantly, it also means more trade investments and prominent retail stocking leading more consumers to buy (see, Livsey, “Wine expert Robert Parker leaves a pointed legacy,” *The Financial Times*, 12/16/2016).²

- In fashion industry new fashion trends are influenced by key opinion leaders and critics. Recommendations of powerful opinion leaders like Anna Wintour of Vogue magazine or Glenda Bailey of Harper’s Bazaar have crucial influence on whether consumers adopt new designs. But such recommendations are also critical because they can convince major retailers to invest in and promote the new design.
- Some art critics influence the reputation of artists as well as the price of their work. For example, according to art historian Kenneth Clark: “In so far as taste can be changed by one man, it was changed by Roger Fry.”³ An English painter and art critic, Roger Fry wrote extremely influential art reviews, capable of making or breaking reputations.
- For technology startups the backing from well known VC firms such Kleiner Perkins or Andreessen Horowitz, not only implies early stage investments, but also influential endorsements affecting their ultimate success. Such endorsements can help not only to attract consumers, but also high quality employees, suppliers, and developers, thereby increasing the likelihood that the product will be successful and will beat out competition:

At the core of these examples is the idea that consumers are uncertain about which

²The critic’s influence on the retail trade is summarized by Max Lalondrelle, fine wine buying director for the London vintner Berry Bros and Rudd: Nobody sells wine like Robert Parker. If he turns around and says 2012 is the worst vintage I have tasted, nobody will buy it, but if he says its the best, everybody will. (see Edgecliffe-Johnson, “Robert Parker American Bacchus,” *The Financial Times*, 12/14/2012).

³Ian Chilvers. “Fry, Roger.” *The Concise Oxford Dictionary of Art and Artists*. 2003.

product will succeed and end up being the best one. Influential recommenders may then have a coordination role by convincing consumers to adopt a product. And central to the coordination role is the existence of feedback effects: A higher *Business Week* ranking helps a school to attract more donations, or the backing from an influential VC can help a technology startup to hire high quality talent. One goal of this paper is to demonstrate the manner in which these feedback effects make the recommendations come true thereby acting as self-fulfilling prophecies. Such feedback effects enable recommendations to have influence even if the recommender were to only have imperfect knowledge about the qualities of the different products in the market.

Specifically, we look at a game with the following characteristics: i) The market consists of N firms each of which may potentially be able to improve quality and produce the best product, ii) $N - M$ firms have obsolete products (which do not match consumer tastes) and only the recommender knows the identity of those firms, iii) Firms compete in quality and prices to win consumers who are uncertain about the firm types, but can listen to a recommendation agent about which firm might be able to produce the best product, iv) the ultimate quality of a firm's product depends upon the number of consumers that are expected to buy the product, v) the recommendation agent's objective is to be accurate which means that the recommended product ends up with the highest quality.

When quality improvement and obsolescence are relevant, the equilibria of the game involve the recommendation agent selecting any one of the M viable firms, all consumers following this recommendation, and only the recommended firm successfully improving quality so that the market indeed observes this firm to be of the best quality. If the recommended firm were to expect that the recommendation would lead all consumers to buy its product, then it will have the resources to invest in developing the highest quality product, i.e., the feedback effect. And consumers who ignore the recommendation will risk ending up with a low quality or obsolescent firm. Recommendations are thus self-fulfilling.

Several implications arise from the basic insight of recommendations as coordination mechanisms with feedback effects. Such recommendations can confer some monopoly power to firms even in markets where consumers preferences are ex-ante fully homogeneous. The recommended firm has monopoly power in the sense that it can charge a price premium and

get the full market demand despite the fact that all the other firms price at marginal costs. Recommendations generate monopoly power even if the recommender has little information about the firms and even if she does not recommend the firm with the highest potential quality.

Moreover, the recommender does not have any incentives to acquire information about potential innovations even when the acquisition costs are small. The recommender can be correct by simply facilitating coordination among consumers around her recommendation, while consumers know that in equilibrium they are purchasing the best product in the market. The resulting equilibrium is thus inefficient since information acquisition is neglected and consumers coordinate around an uninformative recommendation which comes true.

We analyze the problem of a recommender who has to choose between bandwidth (the ability to reach a larger audience) and precise information about firm types. We show that the recommender favors bandwidth over information since with feedback effects greater bandwidth is enough to guarantee that the recommender is accurate, while information precision is not. An associated implication pertains to the case in which the recommender does not know for sure and has an informative but imperfect signal about whether the firm that she recommends is among the firms that can improve and have the best potential quality. Our analysis shows that no matter how precise the signal is, the recommender will ignore her information and instead make a safe recommendation of the firm which has the highest observed quality, rather than the firm which her signal indicates will be able to improve and produce the best quality. Thus even a tiny amount of recommender uncertainty can lead to valuable information being ignored. This result highlights a rational basis for the existence of status quo biases in the development of new products.

In the basic model, the recommender is indifferent between the M equilibria. Another source of status quo bias may arise simply from the lower costs involved in researching and writing a report for the product with higher observed quality, perhaps because more attention has been placed on the product before. Any small cost is enough to break the recommender indifference. Taken as a whole, these results highlight how recommendations with feedback effects lead to inefficient information acquisition and transmission thereby

suppressing successful innovation in many markets.

We also show that the recommender avoids risky innovation products which have a chance of failure even if all consumers were to coordinate around them. Because the recommender cares about accuracy, she will not recommend such products even if the chance of success and quality improvement are arbitrarily high. This incentive for the recommender to play it safe therefore creates substantial inefficiencies in the amount of innovation that is generated, as firms in turn will have the incentive to avoid engaging in innovative strategies if they are associated with even a small amount of risk.

Next we analyze the role of consumer information and heterogeneity on market outcomes. Consider consumer heterogeneity in quality valuations which results in differentiated competition between firms. The recommender can still induce a full coordination equilibrium with all consumers buying the recommended product as long as the spread in the valuations (or the proportion of high valuation consumers) is not too high. Thus recommendations can lead to higher quality innovations being produced in more homogeneous markets.

Finally, we study competition between recommenders. First, we argue that the recommendation market has the characteristics akin to a natural monopoly by showing that it is difficult for even a perfectly informed entrant with lower bandwidth to compete with an uninformed incumbent as the difference in bandwidth between the two makes the incumbent always accurate. Second, we show that even when the market is split between two symmetric recommenders, inefficiencies persist or are even amplified.

Note that the recommendation equilibria in our model are fragile and any firm in the market can be recommended and become successful. What contributes to this fragility? First it is the fact that the quality improvements are assumed to be substantial enough such that any selected firm has the potential to become the best firm as long as the whole market coordinates on it. Second, and perhaps more interesting is the role of the recommender's payoffs being driven solely by the reputation for accuracy. If the recommender cares not only for accuracy, but also for direct monetary payoffs from the recommendation, then this can provide incentives for coordination around a specific (and potentially the most efficient) firm. For example, perhaps contrary to the commonly held view, if firms have

private information about their quality improvement abilities and are able to pay for the recommendation, it is possible that only the most efficient firm will be selected and will end up producing the highest quality. Furthermore, direct payments from the firms may also counter the information acquisition and the risk taking inefficiencies that are associated with recommendations with feedback.

It is important to note that the information acquisition and risk taking inefficiencies created by the recommender exist *precisely* because the recommender has reputational concerns for accuracy. It is because the recommender wants to be correct that she does not take risks or acquire additional information. This suggests that the reputation for accuracy has perverse consequences for recommendations with feedback leading to the suppression of quality improving innovations and welfare loss.

II Related Research

At the broad level this paper is related to literature following Crawford and Sobel (1982) on strategic communication in advice games between an advisor (recommender) who seek to influence receivers who are decision makers when the advisors preferences are inherently different from those of the decision makers. Sobel (1985) introduces reputation effects to advice games. The decision maker is uncertain about the recommender's preferences, so that the recommender's past reports determine his credibility. Our paper is related to the literature where the advisor cares about his reputation for accuracy. Ottaviani & Sorenson (2006) analyze the reporting of private information by an expert with reputational concerns for being perceived as having accurate information and investigate the nature of the information loss. While the recommender in our paper is concerned about accuracy, the communication is aimed at coordinating the market because this helps to build the best quality and to make the recommendation come true.

Another strand in the advice literature concerns reputation for being honest: Morris (2001) shows how such reputational concerns can generate perverse incentives for a good advisor who is honest and who wishes to separate from a bad advisor. If a bad advisor is biased towards a certain message, then a good advisor may avoid sending that message even

when it is accurate, to avoid damaging his reputation.⁴ Thus in this literature the advisor may well send an inaccurate message which strengthens his reputation for incorruptibility, whereas in our model the recommender wants to be accurate and any inefficiency in who is recommended arises precisely from the recommender's message being correct in equilibrium.

The Bayesian persuasion literature originating in Kamenica and Gentzkow (2011) studies the general problem of how a sender can persuade and influence the actions of a rational decision maker by controlling the informational environment even when the sender cannot distort any information that is conveyed. By choosing the signal structure optimally the sender can rationally influence the decision maker's beliefs and induce her preferred action with greater probability while at the same time balancing this with the adverse belief which leaves the decision maker's action unchanged. In this paper the mechanism that makes the recommender's message influential is the coordination of consumer decisions around the selected firm, because in equilibrium due to feedback effects the recommended firm is indeed the best for consumers.

The literature on coordination games (Cooper and John 1988) highlights the role of strategic complementarities in the agents' payoff functions (in the sense of Bulow et. al 1985) as a basis for coordination problems. The feedback effects discussed in our paper allows recommendations to serve a coordination role around any of the sellers. Indeed even if the (multiple) equilibria in our model can be Pareto ranked, the recommender can recommend any firm and still be correct. This fragility due to which any firm can be selected and be ultimately successful underlies several of the information acquisition and risk taking inefficiencies that we highlight in the paper.

Another rationale for information loss comes from the herding and information cascades literature. An important account of herding in investment is by Scharfstein and Stein (1990) who show that managers who care about accuracy can display herd behavior, if after their investment decision the market updates its beliefs based on the success of the

⁴In a related vein, Durbin and Iyer (2009) show that the presence of third-parties can endogenously create conflicts of interests leading to perverse reputational incentives for the recommender and the loss of information transmission. Morgan and Stocken (2003) show how stock analysts who have incentives that are aligned with the investor communicate unfavorable and not favorable information.

investment and also whether the manager's behavior was similar to other managers. The other influential account of information loss comes from the herding and information cascades literature (Banerjee 1992 and Bikhchandani et. al 1992) where agents who observe the actions of others who decided before them, choose to follow while ignoring valuable private information. In our model the recommender's action provides a coordination device which simultaneously gets consumers to herd on one of the firms and the firm to successfully produce the best quality. Further, in contrast to these papers, we study the incentives for the recommender to acquire information and take risks.

The mechanisms for feedback in our model such as R&D investments by firms or the ability to hire better employees generate what in the reduced form is analogous to network externalities as in Katz and Shapiro (1985). The quality improvement that a firm can produce depends upon the number of consumers who are expected to coordinate on the firm even as the consumers willingness to buy the firm's product depends on the quality that will be produced. This paper highlights the role of strategic recommendations with feedback in endogenously generating coordination and network externality like effects that is missing in the literature. Two contributions emerge from this coordination game: First, recommendations with feedback provide a micro foundation for network like effects to endogenously emerge in equilibrium. Second, we are able to identify important and persistent inefficiencies arising from the presence of the recommender which lead to the loss of valuable information and inefficient quality innovation.

Manso (2013) examines feedback effects in the context of credit ratings: A credit rating affects the interest rate, which affects the default decision of the issuer, which in turn influences the rating.⁵ As in this paper, the recommender (credit rating agency) cares about accuracy, while her recommendation (rating) itself can affect the quality of the recommended product. However, we go beyond here by explicitly modeling the coordination problem and investigating how the resulting feedback effects influence the product market competition and the recommenders' incentives. This allows us to investigate the different types of inefficiencies created by the recommender.

⁵For other models of credit ratings with feedback effects, see Boot, Milbourn, Schmeits (2006) and Goldstein and Huang (2019).

III The Model

Consider a market with N firms each of which can potentially develop a product and compete for a unit mass of consumers. Each consumer has unit demand and values products based on their quality. Each firm has a publicly observable initial quality and firm i 's initial quality is denoted by $a_i > 0$. Each firm i can also potentially develop and improve its quality by $\tilde{b}_i x$, where \tilde{b}_i is the firm's ability to improve and x is the fraction of the consumer market that will buy the firm's product. The firms' \tilde{b}_i 's are independently distributed random variables with support $[\underline{b}_i, \bar{b}_i]$ where $\underline{b}_i > 0$.

Once the product is developed all firms have constant marginal costs of producing the product which is normalized to zero. The initial quality a_i can be interpreted as all the existing quality information of Firm i which is publicly observable and known to the market. Note that the firm's quality improvement $\tilde{b}_i x$, implies that the extent of the quality improvement for a firm will be endogenous to demand or the number of consumers who end up buying the firm's product. As will be seen in the analysis, this assumption becomes the basis for the feedback effects arising from recommendations. The underlying mechanisms representing this quality improvement formulation can be illustrated through the following examples:

- *Example 1, R&D Investments:* In technology markets costly R&D investments may be necessary to improve quality and the demand that the firm expects to get determines the amount of R&D for quality improvement that it would be able to invest. Suppose that the firm can improve quality by δ by incurring a R&D cost of $\frac{\delta^2}{2b_i}$. If the firm expects to get a fraction x of consumers to buy and they value the improvement at δ , then the firm will optimally choose to improve quality by $\delta = b_i x$.
- *Example 2, Donors, Investors:* In the business school rankings context better rankings which increases the applications/demand for a school may also motivate greater donations from status conscious donors which provides greater resources for quality improvement investments. For instance, donor giving for firm i can be $b_i x$ if the donor expects the demand to be x , and this when invested produces an equivalent amount of quality improvement units.

- *Example 3, Hiring Talent:* The expectation that a start-up firm will have the best quality and greater demand after the recommendation can lead to higher quality developers or engineers to join leading to better R&D and greater quality improvements.

Next assume that after the firms produce the product quality only some of the firms will end up having viable products while others become obsolete or are a bad match for consumers. Specifically firm i 's quality is given by $\tilde{\theta}_i(a_i + \tilde{b}_i x_i)$ where M firms have viable products ($\tilde{\theta}_i = 1$) and $N - M$ firms will have products that are obsolete ($\tilde{\theta}_i = 0$).

Consumers are uninformed about firms' quality improvement abilities \tilde{b}_i 's, and they only know the distribution of the abilities. So they value any information that resolves this uncertainty. Consumers also do not know which firms will turn out to be obsolete. Next we have the following two assumptions on the parameters a_i and b_i such that,

Assumption 1 *Quality improvement matters:*

$$a_i + \underline{b}_i > a_j \quad \forall i, j \in N \quad (1)$$

Assumption 2 *Obsolescence matters:*

$$\frac{M - 1}{N - 1} < \frac{\min_{i \in N} a_i}{\max_{i \in N} (a_i + E[\tilde{b}_i])} \quad (2)$$

The first assumption implies that any firm can potentially produce the highest quality product in equilibrium if it were to capture the entire market. The second assumption requires that there are enough firms in the market that might end up being obsolete.

The market has a recommender R who can provide product recommendations to consumers. The recommender's objective is to be accurate, which means that the recommender gets payoffs if the recommended product is produced and is of the highest observed quality. At the start of the game the recommender knows the realization of $\tilde{\theta}_i$ for each firm but is uninformed about the \tilde{b}_i 's. Thus while R has information on which firms might become obsolete, she has no better information on the quality improvement abilities of the firms as compared to the market. This sets the stage for the potential coordination role played by the recommender.

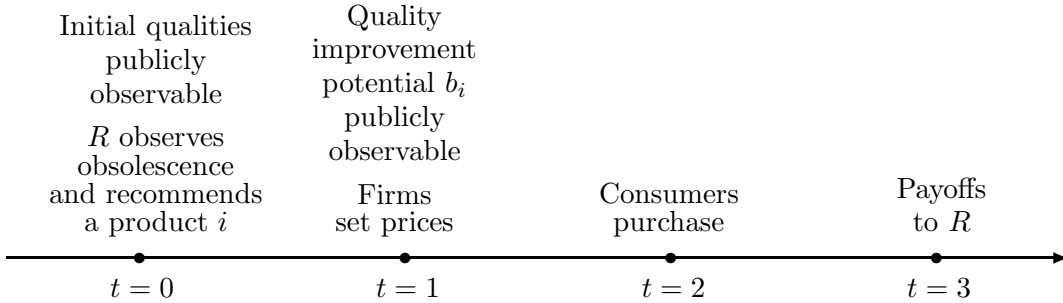


Figure 1: Timing of the Model

The timing of the game is as follows: At time $t = 0$, the recommender makes a recommendation for Firm i . At this time the initial qualities a_i 's are publicly observable, but only the recommender knows (and not consumers) the realization of $\tilde{\theta}_i$'s and which M products will be viable. Next, at $t = 1$, all firms simultaneously choose prices and the quality improvement ability \tilde{b}_i becomes public. At time $t = 2$, consumers make purchasing decisions having observed the prices, the product recommendation, as well as \tilde{b}_i . The assumption that recommendations reveal the quality improvement of the recommended product prior to purchase, while consumers remain uncertain of the qualities of all other products, represent several aspects of markets. For instance, if consumers have search costs of sampling another product, they may eschew search and purchase only based on the expectations of the other firms' qualities.⁶ Another, interpretation is that the recommendation focuses attention on the product, makes it salient, and/or induces word of mouth learning about the quality of the recommended product prior to purchase. Recommendations can make a product salient, not only because it garners attention of consumers, but also other entities such as distribution partners, the press, or other opinion leaders. Finally, at $t = 3$, all qualities are observed and the recommender's payoffs are realized.

Throughout the paper we will focus on pure strategy sub-game perfect equilibria of the game.

IV Self-Fulfilling Recommendations

We now proceed to establish the basic coordination role of recommendation with feedback effects in the Proposition below. The following proposition presents the equilibrium solution:

⁶Indeed, even if the search costs were to be very small and consumers search products sequentially, then in the spirit of Diamond (1971) they would not have the incentive to search additional firms.

Proposition 1 *The game has M equilibria. They involve the recommender R selecting firm $i \in M$, all consumers following the recommendation, firm i charging*

$$p_i = a_i + b_i - \left(\frac{M-1}{N-1} \right) a_j$$

and firm j charging 0, where

$$j \equiv \arg \max_{k \in N \setminus \{i\}} a_k.$$

Proof. Suppose the recommender R were to select some viable firm $i \in M$. In the proposed equilibrium consumers will obtain expected surplus of $\frac{M-1}{N-1}a_j$. No individual consumer will be able to make herself better off by deviating to purchasing from another firm. If a consumer were to ignore the recommendation, she will purchase from the second-best firm at zero price obtaining surplus equal to $\frac{M-1}{N-1}a_j$. The recommended firm i has no incentive to deviate. In equilibrium, firm i profits are p_i . If it increases its price, its profits will be zero as all consumers will move to buying from the second-best firm j . On the other hand, if it lowers prices, profits will be lower since demand will be unaltered. All the other firms will have zero demand at any price higher than 0, and will have negative profit at any price lower than 0. Given the above, R has no incentive to deviate in the proposed equilibrium since her recommendation is accurate due to Assumption 2.

There are no other equilibria. Suppose there was an equilibrium in which a share of consumers ignore the recommendation. These consumers would then have to purchase from some firm $k \neq i$ who may have an obsolete product. Assumption 2 guarantees that even if they pay zero for the product, they would be worse off than following the recommendation because of the risk that they may be buying an obsolete product ■

This proposition establishes the basic result pertaining to the coordination role of the recommender under arguably a modicum of informational requirements. Even if the recommender has no information on the \tilde{b}_i 's, she still has the power to sway the market to adopt any product that is recommended. In the presence of feedback effects if all consumers are expected to follow the recommender's selection, then that firm is indeed able to produce the highest quality. Not only does this make it rational in equilibrium for consumers to buy from the recommended firm, but in turn, R 's recommendation also ends up being accurate. Feedback mechanisms therefore allow the recommender to play a coordination role even in the presence of a little informational advantage. Nevertheless, only one of the

M equilibria will be efficient in the sense that the product with the highest potential quality gets recommended by R .

To clarify the feedback mechanism consider the example of R&D investments made by firms to endogenously choose quality improvements. At $t = 1$ of the game if all firms simultaneously choose their quality improvement investments δ_k and prices by incurring the cost $\frac{\delta_k^2}{2b_k}$, then in the equilibrium the recommended firm i will improve quality by $\delta_i^* = b_i$ and all others will invest zero. The implication is that the recommended firm invests precisely because in equilibrium it is the one that expects to have the resources to do so. This in turn allows the firm to be the best quality firm and thereby generate the resources.

Note that the recommender is indifferent between the M equilibria indicating the fragility in the equilibrium outcomes of this model. This fragility results from the fact that R 's payoff is based on a reputation for accuracy: i.e., it depends purely upon whether she is correct in equilibrium. Any side payment, contractual mechanism, or product preference which makes R 's payoffs dependent upon who is selected might nullify recommender indifference and remove the fragility of the equilibria. Finally, notice also the role of firm competition in this market. Competition a la Bertrand between the recommended and the second best firm implies that the recommended firm i cannot exercise full monopoly power and must price so as to leave consumers with a surplus of $\left(\frac{M-1}{N-1}\right) a_j$ proportional to the quality of the second best firm. However, the recommended firm is able to retain the full surplus from the quality improvement b_i that results from the recommendation. The higher the obsolescence risk, the higher the surplus kept by the recommended firm.

The recommendation equilibria under feedback in this model become self-fulfilling precisely because the coordination facilitated by the recommender prevents the production of any counterfactual information: Even if a number of other firms in the market had better quality improvement abilities compared to the recommended firm their actual observed quality would never be higher. This indicates an information based inefficiency in the development of quality innovations which we investigate in subsequent sections.

V Information Acquisition by the Recommender

The basic setup above establishes the coordination role of the recommender in a world where both the recommender and the consumers do not have information on the quality improvement capabilities of firms. This leads to $M - 1$ of the equilibria described in Proposition 1 to be inefficient in the sense that they lead to any viable firm being selected and ending up being of the highest quality in equilibrium, despite not having the highest quality improvement ability. This naturally leads to the question of the incentives for information acquisition by the recommender and the consumers.

A Costly Information Acquisition

Suppose R can acquire information on the identity of the best potential firm. Specifically, assume that the recommender can choose to incur an acquisition cost c to learn precisely the identity of the best potential firm among those that will have viable products in the market: i.e., the firm that will have the highest possible quality $a_i + b_i$ if it were to sway the market and to get all the demand. The following proposition describes the information acquisition incentive:

Proposition 2 *In equilibrium, the recommender does not acquire information. The equilibria of the game are the same as in Proposition 1.*

Proof. The recommender will still be accurate by recommending any of the M firms. Therefore, she has no incentives to acquire information and incur the cost c . ■

As we know from Proposition 1 the recommender can recommend any one of the firms and still be accurate as that firm is able to improve quality in equilibrium. Thus there is no incentive to indulge in costly information acquisition. At the more general level this highlights an important inefficiency: Recommendation agents do not have incentives to invest in searching for new ideas or innovations in markets with feedback because they can coordinate the market and make their selection come true. This lack of incentive for information acquisition is also driven by the nature of the payoffs of the recommendation agent, in that the recommendation agent cares only about being accurate. To highlight

this point, note that even if R were to be given full information about all the b_i 's for free and is perfectly informed, Proposition 1 will continue to hold and any one of the firms may potentially be selected in equilibrium. As will be seen in the subsequent analysis when the recommender cares either directly or through contractual incentives about the extent of the surplus created, it is possible for both information acquisition incentives as well as efficiency in recommendations to be restored.

B Bandwidth versus Information

One of the strategic trade-offs that recommenders may face is whether to invest costly resources in knowing more about the firms' quality improvement ability or to improve their bandwidth so as to reach more consumers. Suppose that the recommender has limited bandwidth and can only reach a fraction β of consumers. Assume that for all $i \in M$, there exists a $j \in M$ such that

$$a_i + \beta \underline{b}_i < a_j + (1 - \beta) \bar{b}_j. \quad (3)$$

This assumption implies that even if all consumers who are reached by the recommender follow her recommendation, the recommended firm's product quality may be inferior to at least one of the other viable firms in the market. As before we maintain the assumptions in (1) and (2). Suppose that the recommender were to choose between perfect quality improvement information and full bandwidth. How would the recommender choose between the two?

Proposition 3 *The recommender prefers full bandwidth over perfect quality improvement information.*

Proof. Full bandwidth ensures quality as the recommender is accurate in the M equilibria of the game. Acquiring perfect quality improvement information does not guarantee that the recommender will be accurate, since under (3) even if all consumers reached by the recommender follow her recommendation, there could still exist higher quality products in the market. In fact, the $1 - \beta$ consumers not reached by the recommender may avoid purchasing from any firm with positive price since they are uncertain of the quality. Positive prices do not reveal quality since any firm(s) with zero demand can mimic those prices. ■

The recommender would thus be willing to give up information about \tilde{b}_i in exchange for bandwidth. For example, in the case of business school ratings the circulation and hence the reach of the publication is crucial in how much its ratings influence deans as well as readers. Publications care about maximizing the circulation of their business school ratings issue. However, as we show above this preference for increased bandwidth can also mean inefficiencies in the level of innovation with the recommender not necessarily selecting the product with the best quality improvement potential.

VI Status Quo Bias

In this section, we show several reasons for the existence of status quo bias, i.e., the tendency of the recommender to stick to the current market leader who has the highest existing quality (i.e., the firm with the highest $a_i, i \in M$).

A Costly Recommendations

In the basic model, it does not cost the recommender to make a recommendation. However, because making recommendation requires at the very least writing a report, it may be costly for the recommender to do so. Recommendation costs may differ across different products due to the recommender's past experience with products, ease of evaluating products, or due the amount of information that is readily available to evaluate a product.

We assume that the recommender faces private cost c_i to recommend product $i \in N$. Moreover, assume that the status quo product (the one with the highest a_i for $i \in M$) also has the lowest cost. This is natural because recommenders are more likely to have written reports on market leaders in the past or because there is already more attention and information on that product. The presence of even tiny cost differentials can perpetuate the prevalence of status quo options.

Proposition 4 *The only equilibrium of the game involves the recommender issuing a recommendation for the status quo product (the one with the highest initial quality a_i). Firms set prices as in Proposition 1 and all consumers follow the recommendation.*

Proof. In Proposition 1 the recommender was indifferent among playing the M equilibria. Under the assumption of different costs the recommender prefers to select the firm with the lowest cost of writing a report. For similar reasons to the one in Proposition 1 there are no other equilibria of this game. ■

Because of the fragility of the equilibria in Proposition 1 any small cost difference can induce the recommender to recommend a particular product. The recommender is likely to have lower costs to write a report on the status quo product since there is already more attention on that product and the recommender is likely to have written a report on the product before.

B Not All Products Can Improve Quality

In the basic model the assumption $a_i + \underline{b}_i > a_j, \forall i \neq j$ ensured that any firm could produce the highest quality if it were to capture the entire market. This made it safe for R to recommend any one of the N firms. No matter which firm was recommended the recommender would be accurate in equilibrium. We now examine the implications of relaxing this assumption by assuming that $\underline{b}_i = 0$ for all $i \in N$. This relaxes Assumption (1) as $a_i + \underline{b}_i < a_j$ for some $i, j \in N$. There may exist some low ability firms which may not be able to produce the highest quality even if they were to win the entire market.

Proposition 5 *The only equilibrium of this game involves the recommender selecting the viable firm with the highest initial quality a_i , the firm pricing the product at $p_i = a_i + b_i - \left(\frac{M-1}{N-1}\right)a_j$, and consumers following the recommendation.*

Proof. By selecting the viable firm with the highest a_i the recommender is accurate in equilibrium. Firms and consumers also don't have incentives to deviate for the reasons explained in Proposition 1. No other equilibria exist. In equilibrium, consumers must follow the recommendation to avoid purchasing from a firm with unknown quality. If R were to recommend any other viable firm j , there is a risk that she will not be accurate, since it may be the case that $a_j + \tilde{b}_j < a_i$. ■

Recommending any firm which is not the one with the highest a_i implies that there is a chance of selecting a firm j with $a_j + \tilde{b}_j < a_i$, which will not be the best product available in the market even if all consumers were to buy the product. This creates a risk

for the recommender to make an inaccurate recommendation. The recommender can always recommend the firm with the highest a_i and ensure that her recommendation is accurate. Thus there is a rational incentive to stick with the status quo option which has the highest existing quality. This in turn suppresses the supply of quality improvement innovations in the market.

Note that the advisor has a status quo bias even if she learns beforehand that the quality improvement capability of the firm with the highest a_i is very small or non-existent, and even if there are numerous other firms which could improve quality beyond the highest a_i . This suggests that the inefficiency in the equilibrium quality can be quite severe.

VII Risk Taking (In)efficiencies

We have shown that the recommender is able to influence the market in a world where the selection of any firm acts as a coordination device. This is true even when the recommender has no additional knowledge about the individual firms' quality improvement ability. In this section we look at factors which threaten the recommender's ability to be accurate and examine whether and how they affect the efficiency of quality improvement.

A Recommender's Risk Avoidance

We ask how uncertainty about a firm's quality improvement ability affects R 's incentives. Suppose there is an additional firm, $N + 1$, which with probability r is the firm with the highest potential for quality improvement, and with probability $1 - r$ will be dominated by some other firm(s) even if it were to win the entire market. Specifically, suppose with probability r , $N + 1 = \arg \max_{i \in N+1} (a_i + \tilde{b}_i)$, and with probability $1 - r$, let $a_{N+1} + b_{N+1} < a_i$ for some $i \in M$.

Proposition 6 *In equilibrium, the recommender never recommends firm $N + 1$. The $M - 1$ (or M if firm $N + 1$ is obsolete) equilibria are the same as in Proposition 1 with equilibrium prices appropriately adjusting for the additional firm's a_{N+1} .*

Proof. It is easy to see that the equilibria described in Proposition 1 survive in this setting.

Recommending firm $N + 1$ is not an equilibrium action for the recommender as with probability $1 - r$ the recommendation will be inaccurate. R could be accurate with probability 1 by recommending one of the other firms that are not obsolete. For similar reasons to the one in Proposition 1 there are no other equilibria of this game. ■

The recommender refuses to take any risk because she is concerned about her reputation for accuracy. The firm $N + 1$ will never be selected by R even if it is almost certain to innovate and to improve quality ($r \rightarrow 1$) and even if its quality improvement ability is large compared to the other firms. This incentive for the recommender to play it safe is another reason why there may be inefficiencies in the amount of innovation that is generated in the market. This also indicates why it is hard for new firms to establish themselves in the market as the best firm even if they have significant upside potential in producing innovations. Even small amounts of uncertainty about such firms can deter the recommender.

B Firms' Risky Innovation Investments

In the base model, firms are assumed to be endowed with an improvement potential \tilde{b}_i . Here we discuss what happens if firms can make ex-ante investments in risky innovation. First, we assume each firm i can make an investment k in risky innovation which increases product quality by $\delta_H > 1$ with probability p and decreases product quality by $\delta_L < 1$ with probability $1 - p$. The timing is as before, but firms make the risky innovation investment before the recommendation is issued. The recommender observes the investment decision but does not observe the outcome of the risky innovation investment. Given this we have:

Proposition 7 *In equilibrium, firms do not invest in the risky innovation project. The M equilibria of the game are the same as in Proposition 1.*

Proof. Suppose firm i were to invest in the risky innovation as specified above. R upon observing the investment will not recommend firm i as it does not observe the outcome of the investment. Doing so will entail the possibility that the recommended firm does not have the highest quality. Given this firm i does not invest in the risky innovation project. For similar reasons to the one in Proposition 1 there are no other equilibria of this game. ■

Even if the innovation project is really promising (p is close to 1 and $p\delta_H + (1 -$

$p)\delta_L \gg k$), firm i will still pass on the project because the firm knows the recommender is averse to risk. Further, allowing the recommender to acquire costly information about product quality improvement does not affect the results for the same reasons as in section V.

VIII Consumer Heterogeneity

In the basic analysis we have shown the coordination ability of the recommender when the consumers are homogeneous. A homogeneous market precludes the possibility of firm differentiation and presumably allows the recommender to have a greater role in influencing the market outcome. We can therefore ask whether the ability of the recommender to influence coordination might be affected by the presence of heterogeneity in the consumer market. We investigate below the role of heterogeneity in consumer valuations for quality.

Consider the basic model but with consumers who are heterogeneous in their valuation for quality. Specifically, a proportion γ of consumers have higher valuation v_h while the remaining consumers have valuation $v_l < v_h$. Let Υ denote the relative valuation $\frac{v_h}{v_l}$. Assume for simplicity and without loss of generality that the quality improvement abilities of the firms \tilde{b}_i 's are independently distributed random variables with common support $[\underline{b}, \bar{b}]$ where $\underline{b} > 0$. With consumer heterogeneity in valuations there is a possibility for quality differentiation in the market. Might it be that despite consumer heterogeneity the recommender could still sway all consumers to buy the selected product such that it is the only product that can improve quality? Define as before $A_j = \frac{M-1}{N-1}a_j$. In the following proposition we identify the equilibria of this game to address this question:

Proposition 8 *The following are equilibria of this game:*

- i. When $\Upsilon \leq \frac{a_i + b_i - A_j}{\gamma(a_i + \gamma b_i - (A_j + (1-\gamma)\frac{M-1}{N-1}E(\tilde{b})))}$, there is an equilibrium which consists of R selecting any one of the M viable firms, all consumers following the recommendation, the recommended firm i charging $p_i = v_l(a_i + b_i - A_j)$, and the next best firm j charging 0, where $j \equiv \arg \max_{k \in N \setminus \{i\}} a_k$.
- ii. When $\Upsilon > \frac{a_i + b_i - A_j}{\gamma(a_i + \gamma b_i - (A_j + (1-\gamma)\frac{M-1}{N-1}E(\tilde{b})))}$, and $a_i + \gamma b_i - (A_j + (1-\gamma)\frac{M-1}{N-1}E(\tilde{b})) > 0$,

there is an equilibrium which consists of R selecting any one of the M viable firms, the recommended firm i charging $p_i = v_h(a_i + \gamma b_i - (A_j + (1 - \gamma)\frac{M-1}{N-1}E(\tilde{b})))$, the next best firm j charging 0, and only the γ high valuation consumers buying from the recommended firm. Under the conditions stated, the recommended firm still has the highest equilibrium quality.

Proof. Suppose R selects firm $i \in M$. Firm i 's pricing strategy will be to either to induce all consumers to buy, or to charge a (higher) price at which to sell only to the h type consumers. Suppose the firm chooses to induce all consumers to buy by pricing at $p_i = v_l(a_i + b_i - A_j)$. In the equilibrium proposed in part (i) of the proposition, no consumer who expects firm i 's equilibrium quality to be $a_i + b_i$ before purchase can make herself better off: A type l consumer will get a surplus of $v_l A_j$ and is no better off if she buys from the second-best firm j , while a type h consumer will only get $v_h A_j$ from the second-best firm instead of $v_l A_j + (v_h - v_l)(a_i + b_i)$. All the other firms in equilibrium also cannot improve their payoffs by charging a different price than zero. This consumer behavior and the prices charged by all firms are consistent with the firm i 's quality being $a_i + b_i$, and the second-best firm having expected quality A_j . Finally, R also has no incentive to deviate as the recommendation is correct.

Next, consider the alternative pricing strategy where firm i prices to sell to only the h type consumers. In the equilibrium proposed in part (ii), the price $p_i = v_h(a_i + \gamma b_i - (A_j + (1 - \gamma)\frac{M-1}{N-1}E(\tilde{b})))$, is the highest price it can charge such that the incentive compatibility constraints of both types of consumers are satisfied (i.e., type h prefers to buy from firm i while type l prefers to buy from the second-best firm). Suppose that $a_i + \gamma b_i - (A_j + (1 - \gamma)\frac{M-1}{N-1}E(\tilde{b})) > 0$, so that this price is positive. No consumer who is atomistic can change their payoffs from deviating, given the expected equilibrium qualities. The consumer behavior and prices charged by all firms are also consistent with the firm i 's quality being $a_i + b_i \gamma$ and the second best firm's expected quality being $A_j + (1 - \gamma)\frac{M-1}{N-1}E(\tilde{b})$. All other firms including the second-best firm also cannot improve their payoffs by charging a different price than zero. Finally, R is correct under the conditions of the equilibrium as firm i will have the highest equilibrium quality.

To complete the proof consider firm i 's payoffs when choosing the alternative pricing strategies: The firm will choose $p_i = v_l(a_i + b_i - A_j)$, if $v_l(a_i + b_i - A_j) \geq \gamma v_h(a_i + \gamma b_i - (A_j + (1 - \gamma)\frac{M-1}{N-1}E(\tilde{b})))$. This leads to the condition $\Upsilon \leq \frac{a_i + b_i - A_j}{\gamma(a_i + \gamma b_i - (A_j + (1 - \gamma)\frac{M-1}{N-1}E(\tilde{b}))}$. Alternatively it will choose $p_i = v_h(a_i + \gamma b_i - (A_j + (1 - \gamma)\frac{M-1}{N-1}E(\tilde{b})))$. ■

A full coordination equilibrium exists when $\Upsilon \leq \frac{a_i + b_i - A_j}{\gamma(a_i + \gamma b_i - (A_j + (1 - \gamma)\frac{M-1}{N-1}E(\tilde{b}))}$. There

are M such equilibria in which R selects a firm i which charges $p_i = v_l(a_i + b_i - A_j)$, and sells to the entire market. When the dispersion in the quality valuations of the consumers are not too high, it is optimal for the recommended firm to price in a manner so as to induce all consumers to coordinate around purchasing from the firm. However, when the quality valuations are sufficiently heterogeneous, then the recommended firm finds it optimal to focus on coordinating the high valuation consumers. The firm charges $p_i = v_h(a_i + \gamma b_i - (A_j + (1-\gamma)\frac{M-1}{N-1}E(\tilde{b})))$, and sells to all the high valuation consumers, while the low valuation consumers find it optimal to buy from the second best firm. With greater dispersion in the quality valuations the recommender might be able to effect only partial coordination in the consumer market. Nevertheless, the recommended quality can still be the highest quality in the market.

It is also useful to note how recommendations and the quality that is produced respond to market heterogeneity. Recommendations lead to higher quality innovations being produced in more homogeneous markets. It is precisely in these markets that a full coordination equilibrium where all consumers buy the recommended product is possible. Further, full coordination is facilitated as long as there are not too many high valuation consumers.

Consumer Information and Price Discrimination: Consider now that the firms have information about the consumer types and are therefore able to (third degree) price discriminate and target prices to each consumer segment.

Proposition 9 *When firms know the consumer types and are able to price discriminate, then the equilibrium always involves full coordination: R recommends firm i and all consumers buy from the recommended firm which is able to produce the highest quality $a_i + \tilde{b}_i$.*

When firms have information on the consumer quality valuations they can use it to price discriminate. Hence the recommended firm can suitably lower the price for the low valuation consumer group so as to induce them to buy, without lowering the price for the high quality valuation consumers. Therefore with price discrimination the resulting equilibrium is one with full coordination. Thus better availability of consumer information which enables firms to price discriminate also makes recommendations with feedback effects

more successful. Furthermore, the ability to price discriminate also leads to the highest quality innovations being produced.

IX Competition Between Recommenders

In the basic analysis there is only one recommender. In this section, we argue that in the presence of feedback there are conditions when recommendation markets can behave akin to natural monopolies and therefore the model with a single recommender is not far-fetched. Moreover, we show that if we allow for competition between recommenders, it can even reinforce the inefficiencies highlighted in the previous sections.

With two recommenders, each recommender is concerned about accuracy, but also values being more accurate than the other recommender. If both recommenders are accurate they split the payoff V according to the number of consumers that follow their recommendation. If only one recommender is accurate she gets the full payoff V .

A Recommendation Markets as Natural Monopolies

Consider two recommenders: an incumbent and an entrant. The incumbent recommender is uninformed about the \tilde{b}_i 's but has full bandwidth, while the entrant is perfectly informed but has small enough bandwidth α such that

$$a_i + \underline{b}_i(1 - \alpha) > a_j + \bar{b}_j\alpha \quad \forall i, j \in N. \quad (4)$$

This means that even if all the α consumers reached by the entrant recommender were to follow her recommendation j that would not be enough to make its product better than another viable product i that the remaining $1 - \alpha$ consumers purchase. The bandwidth disadvantage of the entrant can be interpreted as the presence of cost disadvantage for a second player in the industry to acquire bandwidth.

Proposition 10 :

i) If the entrant recommender moves first followed by the incumbent, then there are M equilibria of the game. The entrant never recommends the best product and the incumbent always recommends the best product.

ii) If the incumbent recommender moves first followed by the entrant, then there are M equilibria of the game. The entrant ignores her information and follows the incumbent's recommendation.

Proof. If the incumbent moves second, she will recommend any non-obsolete product j that is different from the entrant recommendation i . This implies that the incumbent does not need to split the payoff V with the entrant. Assumption (4) guarantees that if at least $1 - \alpha$ consumers will follow the incumbent then product j will indeed be the best product in the market. The $1 - \alpha$ consumers who do not listen to the entrant will stick to the incumbent recommendation because otherwise will run the risk of selecting an obsolete product. We note that consumers who do not listen to the entrant cannot learn the entrant recommendation based only on prices as other firms with zero demand would charge the same price as firm i to try to attract these consumers. If the incumbent moves first, for the same reason as above, her recommendation is guaranteed to be the best product in the market. Therefore, to earn positive payoff, the entrant will ignore her information and follow the incumbent's recommendation. ■

An incumbent with a bandwidth advantage always has the ability to make recommendations that come true in equilibrium, despite the fact that the entrant has perfect information. The effect of recommendation timing on the equilibrium outcomes is worth elaborating: Indeed, if the incumbent is the follower in making the recommendation, then the entrant will never be able to recommend a product that will end up being the best. Given that each firm strictly prefers to be the sole firm that is correct, the incumbent will always have the incentive to recommend a firm that is different from the entrant's recommendation, which given the incumbent's bandwidth advantage assures that the entrant's selection never comes true. The incumbent if it had a choice would always like to go after the entrant and ensure that the entrant's recommendation is wrong. Thus in markets where acquiring superior bandwidth is difficult for an entrant, the incumbent recommender may be able to sustain monopoly recommendation rents.

B Competition and Information Acquisition

What if there are markets in which two large recommenders coexist and compete. Suppose that each of the recommenders are symmetric and can reach half of the consumers. The game unfolds as follows. At each point in continuous time, recommenders maintain a recommendation, which is observable to the other recommender. At any point in time, each recommender can pay an acquisition cost c to learn all the $\{\tilde{b}_i\}_{i \in N}$. It costs $\kappa > 0$ for a recommender to switch recommendations. At some Poisson arrival time \tilde{t} with intensity λ , firms set prices, consumers observe their recommendation and make a purchasing choice.

Proposition 11 *In equilibrium, recommenders do not acquire information and coordinate on any of the equilibria of Proposition 1.*

Proof. A recommender is not willing to acquire information if the other recommender is mimicking her recommendation. Since they will have to share the payoff, there is no benefit to acquiring information just as in Proposition 2. If one recommender starts recommending firm i the other has no incentives to deviate to recommending another firm without additional information because the other recommender will issue a similar recommendation. The small cost κ prevents recommenders from bouncing strategies around in high frequency. ■

Therefore, competition between recommenders does not solve the information acquisition inefficiency in markets with feedback effects.

C Competition and Status Quo Bias

Suppose each of the two recommenders can reach half of the consumers. The two recommenders simultaneously issue recommendations. For simplicity, we assume that the distribution of \tilde{b}_i is binary $\{\underline{b}_i, \bar{b}_i\}$ with probability 1/2 for each outcome.

Proposition 12 *The only equilibrium involves both recommenders issuing a recommendation for firm $i \in M$ that has the highest $a_i + \frac{1}{2}E(b_i)$. Consumers follow the recommendation and the price charged by firm i is the same as in Proposition 1.*

Proof. Consumers will follow their associated recommender to avoid purchasing an obsolete product. If a recommender issues a recommendation for any product different than product i , the other

recommender has an incentive to recommend product i since her probability of being accurate and getting the whole surplus exceeds one half. ■

Therefore, competition can actually create a form of status quo bias, since the only equilibrium that survives is the one associated with the highest $a_i + \frac{1}{2}E(b_i)$, which puts more weight on initial quality than on the quality improvement. In this situation, a monopolist recommender would not have such bias, and could in equilibrium recommend the product with the highest $a_i + E[b_i]$.

X Recommender Payoffs

The analysis up until now has shown how the recommendations with feedback effects affect market outcomes when R 's payoffs come from maintaining a reputation for accuracy. Focusing on reputational concerns helps us to highlight some of the important types inefficiencies including inefficient risk taking, status quo bias, and inefficient information acquisition. The reputation for accuracy also resulted in the fragility of the recommendation equilibria, in that R could recommend any of the M firms which could then potentially produce the best product. In this section, we investigate the implications for the recommendation equilibria when the recommender receives pecuniary payoffs resulting from the recommendation.

In some recommendation markets there is a debate between consumer- versus firm-pay forms of organization. Some recommenders charge fees to consumers for their recommendations. Many of them explicitly state that they do not receive money from firms that they recommend. Other recommenders choose to receive monetary payments from firms (either through advertising revenues or product endorsement fees). We investigate the implications of these different organizational forms in the context of our model.

We first examine monetary payments from consumers to the recommender. Suppose that as in section A the recommender can acquire information about $\{\tilde{b}_i\}_{i \in N}$ at a cost c . Consumers pay a fee to access the recommendation. The surplus to the consumers in the equilibrium derived in Proposition 1 when the recommender selects firm i is $\left(\frac{M-1}{N-1}a_j\right)$ where $j \equiv \arg \max_{k \in N \setminus \{i\}} a_k$. Information about the potential improvement \tilde{b}_i has no material impact on the surplus of the consumers. Therefore, if the recommender and the consumers were

to agree on a mechanism with transfers it would involve no information acquisition by the recommender. There would be a form of bias against the status quo, since consumers could only benefit from the recommender selecting a firm different from the one with the highest a_i as this maximizes their surplus.

Next consider monetary payments from firms to the recommender. One might expect that direct payments from firms to the recommender could induce recommendation bias. Indeed as highlighted in Proposition 1, the recommender is indifferent between the M equilibria. Therefore, any small bribes from a viable firm could presumably win the recommendation creating inefficiencies. In a number of important contexts payments from firms for recommendations are viewed as controversial and inefficient by competition authorities and policy experts. Recently the U.K.’s Competition and Markets Authority launched an investigation into leading social media “influencers” and whether they (and the platforms) were clearly informing followers if they were being paid for endorsements. As another example, post the 2007-09 financial crises policy experts have pointed out the perverse incentives created by the issuer-pay model in which rating agencies are paid by the security issuers whose securities they rate.⁷

Against this backdrop our model suggests that in markets with feedback and coordination effects, payments from firms may have efficiency improving properties. Specifically we argue that an open auction for the recommendation in which all firms participate can restore efficiency. Assume that each firm knows its type $\{\tilde{\theta}_i, \tilde{b}_i\}$. The recommender runs a second-price auction, i.e., the firm with the highest bid pays the recommender the second highest bid and gets the recommendation. In equilibrium, firms will be willing to bid their profit $\theta_i(a_i + b_i - \left(\frac{M-1}{N-1}\right) a_j)$ to get the recommendation. The recommender selects the firms with the highest $\theta_i(a_i + b_i - \left(\frac{M-1}{N-1}\right) a_j)$, taking into account information about the realization of \tilde{b}_i but at the same time introducing status quo bias because $\left(\frac{M-1}{N-1}\right) a_j$ is subtracted from the bids. In other words, there is some incentive for the recommender to favor the highest initial quality firm especially if the next highest initial quality is sufficiently lower.

⁷See Barney Thompson and Aliya Ram, “U.K. probe of “influencers” extended to social media platforms,” *Financial Times*, 01/23/2019. For a critique of the credit-ratings system by a leading expert see Alan Blinder, “A better way to run rating agencies,” *The Wall Street Journal*, 04/17/2014.

If firms do not know their type b_i , the recommender could pay c to learn all the b_i s as in subsection A. Suppose that there is a share contract and the recommender gets α share of the firm profits. For high enough α the recommender can be induced to collect information and select the firm with the highest $\theta_i(a_i + b_i - \left(\frac{M-1}{N-1}\right) a_j)$ alleviating the inefficiency but will still be subject to some status quo bias. It is only if the recommender was able to collect money from firms and consumers in the appropriate proportions that the efficient equilibrium with the highest $\theta_i(a_i + b_i)$ would be selected.

XI Conclusion

The incentives of firms to produce quality improving innovations are often subject to coordination problems. In markets ranging from the ones for technology startups, fashion products, new wine vintages, and business education, there may be non-trivial uncertainty about whether a product will succeed and be the best in the market. The essence of the feedback mechanism highlighted in this paper is that the firm's ability to successfully produce the best quality crucially depends upon whether a significant enough number of consumers are willing to buy it in the first place because this endogenously allows for the resources needed to produce higher quality. And unless the buyers are convinced, they have no reason to adopt the product. Recommendation agents have a strategic role in resolving these types of coordination problems.

Our analysis demonstrates that in the presence of feedback effects the recommendations can have a self-fulfilling property. Any selected firm can potentially be successful in improving its quality to become the best firm in the industry. This may be the case even if the recommender does not have much information on the quality improvement capabilities of the firms. More importantly we have shown that a recommendation agent can end up creating significant inefficiencies in the ultimate quality of the innovation that is produced, precisely because their payoffs are determined by a reputation for accuracy. The recommender will not acquire information even if the acquisition costs are minimal and even when she has significant uncertainty about the firms. The self-fulfilling nature recommendations under feedback mutes the incentive to acquire valuable information. The recommender is

also subject to a rational status quo bias in recommendations: even a small amount of uncertainty about the quality improvement ability of the firms leads the recommender to recommend the firm with the highest existing quality, rather than the firm which can produce the highest total quality. Along the same vein, risky products are never selected by the recommender because going with the safe option guarantees accuracy. Risky innovation projects are also not implemented by firms even if they were to be significantly promising because firms anticipate that the recommenders are averse to risk.

Competition between recommenders is unlikely to mitigate these inefficiencies. Recommendation markets with feedback can behave like a natural monopoly in the sense that it is difficult for an entrants to enter and compete with an incumbent who has bandwidth advantage. Even in markets with competition between symmetric recommenders, the inefficiencies remain or may even be exacerbated. Monetary transfers from consumers and firms can potentially mitigate inefficiency. In contrast to the prevailing view that direct payments (bribes) from firms to the recommender may induce bias, we find that a firm-pay model may restore some efficiency (and is superior to a consumer-pay model), at least when firms are the ones keeping the surplus generated from quality improvement, as a well-designed mechanism may transfer some of this surplus to the recommender.

References

- Banerjee A., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* 107(3) (1992) pp. 797-817.
- Bikchandani S., D. Hirshleifer, and I. Welch "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," " *Journal of Political Economy* 100(5) (1992) pp. 992-1026.
- Boot, A., T. Milbourn, A. Schmeits "Credit ratings as coordination mechanisms," *Review of Financial Studies* 19 (2006), 81118.
- Bulow, J., J. Geanakoplos, and P. Klemperer "Multimarket Oligopoly: Strategic Substitutes and Complements," " *Journal of Political Economy* 93(3) (1985) pp. 488-511.
- Cooper, R. and A. John. "Coordinating Failures in Keynesian Models," *Quarterly Journal of Economics* 103(3) (1988) pp. 441-463.
- Crawford, V. and J. Sobel. "Strategic Information Transmission," *Econometrica* 50(6) (1982) pp. 1431-1451.
- Diamond P.A., "A Model of Price Adjustment," *Journal of Economic Theory* 3(2) (1971) pp. 156-168.
- Durbin E., and G. Iyer. "Corruptible Advice," *American Economic Journal: Microeconomics* 1(2) (2009) pp. 220-242.
- Espeland W.N., and M. Sauder. "Rankings and Reactivity: How Public Measures Recreate Social Worlds," *American Journal of Sociology* 113(1) (2007) pp. 1-40.
- Goldstein, I., and C. Huang. "Credit Rating Inflation and Firms' Investments," *Working Paper* (2019).
- Kamenica, E. and M. Gentzkow. "Bayesian Persuasion," *American Economic Review* 101(6) (2011) pp. 2590-2615.

- Katz M.L., and C. Shapiro. "Network Externalities, Competition and Compatability, " *American Economic Review* 75(3) (1985) pp. 424-440.
- Manso, G. "Feedback Effects of Credit Ratings," *Journal of Financial Economics* 109 (2013) pp. 535-548.
- Morris, S. "Political Correctness," *Journal of Political Economy* 109(2) (2001) pp. 231-265.
- Morgan J. and P. Stocken "An Analysis of Stock Recommendations" *Rand Journal of Economics*, 34(1) 2003 pp. 183-203.
- Ottaviani, M. and P. N. Sorenson. "Reputational Cheap Talk" *Rand Journal of Economics*, 37, 2006, pp. 155-175.
- Scharfstein, D. and J. Stein. "Herd Behavior and Investment," *American Economic Review* 80(3) (1990) pp. 465-479.
- Sobel, J. "A Theory of Credibility," *Review of Economic Studies* 52(4) (1985) pp. 557-573.