A Model of Recommended Retail Prices*

Dmitry Lubensky

Kelley School of Business
Indiana University

December 16, 2013

Abstract

Consumers rely on a manufacturer's recommended price to help determine whether to accept a retailer's price or continue to search. This paper demonstrates that doing so can be rational even if the manufacturer's price recommendation is cheap talk. By incentivizing search, a manufacturer trades off reducing double marginalization and losing consumers to competitors. When the manufacturer's cost is low he induces low retail prices and benefits when consumers search more. When the manufacturer's cost is high he induces high retail prices and benefits when consumers search less. Since consumers prefer to search more when lower prices are available, their incentives are aligned with the manufacturer's and this allows informative cheap communication. Aside from costs, the manufacturer can inform consumers of other market parameters such as product quality.

Keywords: consumer search, sequential search, search with uncertainty, manufacturer suggested retail prices, vertical markets, signaling and cheap talk

*Thanks to Josh Cherry, Kai-Uwe Kühn, Francine Lafontaine, Stephan Lauermann, Doug Smith, and Mike Stevens for insightful feedback on the early version of this paper. Thanks also to Mike Baye, Rick Harbaugh, John Maxwell, Jeff Prince, Michael Rauh, Eric Rasmusen, Collin Raymond, Babur de los Santos, Matthijs Wildenbeest and the participants of the 2013 Consumer Search and Switching Cost Workshop in Moscow for helpful comments on this version.

1 Introduction

Manufacturers use non-binding recommended retail prices in markets ranging from grocery products to big ticket items such as electronics, appliances, and cars. These recommendations, which come in a variety of forms such as list prices, manufacturer suggested retail prices (MSRPs), sticker prices, etc., are visibly printed on product packaging and often promoted by the manufacturer through costly advertising. The existence of a link between price recommendations and market outcomes has both been shown empirically (e.g. Faber and Janssen (2008), De los Santos et. al. (2013)) and also implicitly assumed in the myriad studies that use recommendations as a proxy for transaction prices (e.g. Berry, Levinsohn, and Pakes (1995)). There is also anecdotal evidence that recommendations can directly affect the decisions of market participants, for example, consumers often expect a discount off the MSRP when buying a new car and strategic dealers take this into account as they set prices.[1] However, despite the evidence that price recommendations affect behavior, our understanding of how they do so is quite limited. Due to the fact that price recommendations are non-binding, the mechanism by which they have an impact and the motives of the manufacturer in making these recommendations are still not well understood.

A common explanation is that recommendations act as price ceilings. This story is compelling because most products sell at or below MSRP and because the manufacturer's rationale for imposing a price ceiling in order to reduce double marginalization is well established (Telser (1960), Mathewson and Winter (1984)). Yet this explanation of recommendations is incomplete. First, price recommendations are not binding, at least in name, thus it is not clear why a manufacturer would make a recommendation instead of imposing a price ceiling directly. In addition, recommendations often do not bind in practice, for instance most cars sell strictly below MSRP but very few sell at MSRP. Lastly, manufacturers publicize their recommendations and an explanation of recommendations as explicit price ceilings ignores the potential role played by consumers.

This paper presents an alternative explanation in which price recommendations affect consumer search. Uncertain of market conditions, consumers do not know the distribution of retail prices and sequentially visit retailers. A recommendation can inform them of the distribution of prices and thereby help determine whether to accept their current offer or to continue to shop. Retailers in turn anticipate consumers' reactions to the recommendation and adjust their prices accordingly. Even though by this mechanism they are non-binding, price recommendations have an impact on consumer and retailer behavior, and thus affect real market outcomes such as prices and sales.

[1] While the majority of vehicles sell for prices strictly below MSRP there have been a few notable exceptions such as the Toyota Prius in which prices exceeded recommendations.

Consumers benefit from MSRPs by making more informed search decisions, however the incentive for a manufacturer to provide this information is not immediately clear. I demonstrate that in affecting consumer search the manufacturer faces a tradeoff. When he signals that retail prices are low and encourages consumers to be selective, retailers are pressured to reduce markups which curbs double marginalization. At the same time, encouraging consumers to search makes it more likely that they purchase a competing product. Conversely, when the manufacturer signals that prices are high he increases double marginalization but also lowers the likelihood of losing searching consumers to competitors. Whether the manufacturer prefers to induce more or less search depends on the relative strength of these two countervailing effects.

In order for price recommendations to affect search, the manufacturer must be willing to make recommendations and must be able to convey information credibly. The latter is particularly salient because by their nature recommendations are not falsifiable statements of fact and in this sense are unverifiable cheap talk.[2] If the manufacturer has incentive to mislead consumers, they in turn would rationally ignore recommendations and no information could be transmitted. In order for non-binding non-falsifiable recommendations to have an effect, the manufacturer must prefer to report truthfully.

The paper's main result is that with respect to search, the interests of consumers and the manufacturer can be aligned, which allows for cheap communication. In the model, the downstream market is comprised of retailers exclusive to the manufacturer and other sellers that carry competing products. The manufacturer's marginal cost is uncertain and has a low and a high realization. When the cost is low, the manufacturer charges a low wholesale price which induces low retail prices, thereby making consumers more likely to accept the manufacturer's product than a competitor's. In this situation, inducing consumers to hold out for a good deal increases the probability that they buy the manufacturer's product. Meanwhile, given that retail prices are low, it also benefits consumers to be more selective during search which means their incentives are aligned with the manufacturer's. Similarly, when the manufacturer's cost is high he induces high retail prices and offers consumers a lower utility on average than his competitors. More search now increases the odds that consumers buy from competitors, and the manufacturer prefers consumers to settle for lower utilities. Again, consumers are better off being less selective when the manufacturer's costs are high, thus the interests of the two are aligned when costs are high. Since the two parties agree on whether more or less search is desirable in either state, the manufacturer can credibly convey information to consumers using cheap talk.

[2] With other types of communication, for example printing on the cover of a book that it is on the New York Times Best Seller list, the manufacturer is compelled to disclose truthfully since the statement is verifiable and punishable by law if false. However, price recommendations are not easily falsifiable – the manufacturer can recommend a high price but if retailers choose to ignore this recommendation it is difficult to hold him liable.

In the extensive theoretical literature on vertical price restraints, recommendations have received little explicit attention and instead have been lumped in as an instrument of resale price maintenance. However, there has been a recent renewed interest in the topic and several recent papers have explored the mechanism behind the effects of these non-binding recommendations. Buehler and Gärtner (2013) shows that MSRPs can be used by a manufacturer to convey demand or cost information to a retailer. The manufacturer and retailer are engaged in a repeated game and there exists an equilibrium in which both aim to maximize joint profits. Since in this equilibrium incentives are aligned, the manufacturer can credibly inform the retailer using cheap talk. In this framework, printing MSRPs on product packaging can be a way to coordinate with several retailers at once. However, actively publicized MSRPs are also likely to affect the behavior of consumers, and the authors abstract from this effect.

A different approach is taken by Puppe and Rosenkranz (2011), that builds on the theory in Thaler (1985) in which a price recommendation provides a behavioral reference point. Consumers' demand is kinked above the MSRP so that the recommendation acts as a defacto price ceiling. My paper can be thought of as providing a rational foundation for the reference point theory. I explicitly model how consumers incorporate the MSRP and their knowledge of the market to form price expectations, and derive their demand as the solution to the problem of optimal search. A consumer that refuses to accept a price above a certain level does so not only because it feels like a bad deal, but because it actually is a bad deal.[3]

The role of MSRPs as information for consumers contributes to the policy debate of whether manufacturer recommendations ought to be regulated. When recommendations are banned, consumers remain uncertain about the distribution of prices and likely pursue a strategy in which they search an intermediate amount – more than they would if they knew prices were high and less than they would if they knew prices were low. This harms the manufacturer for whom a move to an intermediate amount of search reduces profit for both high and low cost realizations, which supports the fact that manufacturers voluntarily use recommendations when unregulated. Consumers are also harmed by virtue of making less informed search decisions. However, the net effect on consumers is theoretically ambiguous since using an intermediate search strategy can result in either higher or lower average prices.

The present framework for modeling non-binding price recommendations can also be applied in markets in which producers sell directly. An example is residential housing, in which studies such as Horowitz (1992) have empirically demonstrated the relationship between list

[3]Armstrong and Chen (2012) provides an alternative search rationalization of reference prices. Sellers commit to revealing past transaction prices which consumers use as a reference to infer about current prices charged by competitors.

and transaction prices, but have not established the theoretical foundations for why this relationship should exist. While Horowitz conjectures that list prices help inform consumers of a seller's reservation price, he does not show that such information can be credibly communicated in an equilibrium. By contrast, theoretical models of list prices in this setting rely on them being more than just information. For example, Gill and Thanassoulis (2010) explicitly characterize the effect of list prices on industry competition, but in doing so assume that list prices act as binding price ceilings. I propose a mechanism by which price communications are truly non-binding, and yet have a real impact by providing consumers with information that affects their search.

This paper also provides several methodological contributions. First, I develop a tractable model of a vertical market with sequential search and endogenous prices. By using the Wolinsky (1986) random utility framework I ensure that equilibria are in pure strategies and easily amenable to comparative statics. The closest work in this area is Janssen and Shelegia (2013), which instead uses the Stahl (1989) framework to obtain equilibria in mixed strategies. Janssen and Shelegia's central result is that double marginalization is exacerbated in a setting in which consumers search, and this result also obtains quite naturally in the present model. The papers diverge in that I focus on the manufacturer communicating the realization of uncertain market conditions while in Janssen and Shelegia there is no aggregate uncertainty and thus nothing to communicate. Second, I derive a simple and intuitive expression for the demand faced by an upstream manufacturer in a search environment, drawing a connection between models of search and contests. Finally, the paper contributes to the literature on search with aggregate uncertainty. Models in this literature either have sequential search but only two firms (Benabou and Gertner (1993)) or a larger number of firms and non-sequential search (Yang and Ye (2008), Tappata (2009)). The issue is that sequential search with many firms potentially allows for equilibria where consumers follow non-stationary strategies. I provide a model in which search is sequential and there is a continuum of potential sellers, yet search strategies are stationary because information is communicated credibly using cheap talk. Hence, I show that in principle sequential search among a large number of sellers can be modeled tractably.

The rest of this paper proceeds with Section 2 which presents the model. Section 3 characterizes the equilibrium when consumers are informed, followed by Section 4 which shows the result that cheap signaling can credibly convey information. Section 5 examines the effects of banning MSRPs and draws a comparison between recommendations and resale price maintenance. Then Section 6 concludes.

2 Model

A manufacturer of a product has a constant but uncertain marginal cost, with a low realization c_L and a high realization c_H, both equally likely. The manufacturer sells the product through a collection of N exclusive retailers, which face no additional costs. The downstream market is comprised of these N retailers and also M sellers of other competing products. I will refer to sellers of the manufacturer's product as retailers and to all other sellers as competitors.

There is a measure one of consumers with unit demand. A consumer faces the set of $M+N$ sellers, which she may visit sequentially at a cost of $s > 0$ per observation. A consumer that has visited k sellers and buys from seller j obtains a payoff

$$u_j = v_j - sk,$$

in which v_j is a consumer's value at seller j. If seller j is a retailer, the value is

$$v_j = \eta_j - p_j,$$

the difference between a random utility shock η_j drawn from F and the retailer's price p_j. If seller j is a competitor, the value v_j is drawn from distribution G. Let $[\underline{\eta}, \bar{\eta}]$ be the support of F and $[\underline{v}, \bar{v}]$ be the support of G, with $\underline{\eta} \geq 0$ and $\underline{v} \geq 0$. Assume both distributions admit continuous densities f and g, respectively, and assume that $f(\cdot)$ is log-concave.

The game proceeds as follows. First, the manufacturer draws and privately observes production cost $c \in \{c_L, c_H\}$ and chooses a price recommendation $\sigma \in \{\sigma_L, \sigma_H\}$ and a wholesale price w. Then, every retailer i observes σ and w and sets retail price p_i. With retail prices fixed, each consumer observes the manufacturer's recommendation σ and decides whether to begin searching. If she initially decides not to search, she exits and receives a utility of zero. Otherwise, she pays cost s and visits a randomly selected seller, where she observes p_i and η_i if it is a retailer i or v_j if it is a competitor j. A consumer can accept the offer, exit and receive a total payoff of $-s$, or pay cost s again and visit a different randomly selected seller.[4] During search a consumer can accept any previous offer at no additional cost, and the process continues until she accepts an offer or exits. The search strategy is denoted as $A(v, \sigma)$, in which v is a vector of previously observed offers. The timing is summarized in Figure 1.

The expected payoff of a retailer i when setting price p_i is

$$\pi(p_i \mid w, \sigma) = (p_i - w)q(p_i \mid w, \sigma),$$

[4]The search process is without replacement.

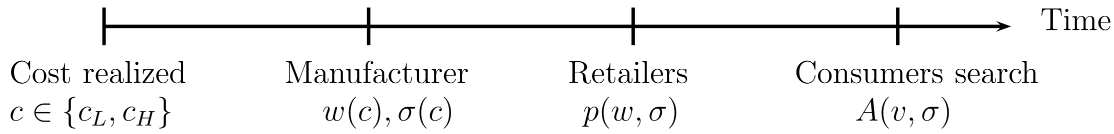

Figure 1: Model Timing

in which $q(p_i \,|w, \sigma)$ is the expected demand faced by retailer i in equilibrium. To form this expectation the retailer can use recommendation σ to anticipate consumers' search behavior and both σ and wholesale price w to anticipate other retailers' prices. The function $q(p_i \,|w, \sigma)$ is an equilibrium object and is derived in the ensuing analysis. The manufacturer's expected payoff is

$$\Pi(w, \sigma \,|c) = (w - c)Q(w, \sigma),$$

in which $Q(w, \sigma)$ is the manufacturer's expectation of the sum of the sales by all his retailers, another object derived explicitly in equilibrium.

The solution concept is symmetric sequential equilibrium, namely a collection of strategies for the manufacturer $(w(c), \sigma(c))$ and retailers $p(w, \sigma)$ and a search policy for consumers $A(v, \sigma)$, all satisfying sequential rationality and beliefs that obey Bayes rule whenever possible.

The present environment borrows from common constructions in the literatures on vertical relationships and search markets. The downstream game is analogous to the random utility search model in Wolinsky (1986) and Anderson and Renault (1999), and the uniform wholesale price is a contract that, while suboptimal, is frequently observed and widely analyzed, starting with Spengler (1950). There are several additional modeling assumptions which are important for the main result and I briefly elaborate on them here.

First, the use of distribution G abstracts from modeling the behavior of competing firms. It will be shown that when signaling his cost, the manufacturer is only concerned with how the distribution of utilities offered by his retailers compares with that offered by competitors; the mechanism by which the latter is determined is not relevant for his signaling decision. Consequently, the main result applies in a variety of market structures, including but not limited to a symmetric setting in which each downstream seller belongs exclusively to one of several competing manufacturers, all facing the decision problem described above. The key assumption is that G is exogenous to the manufacturer, so that he cannot affect the distribution of outside options through his actions. In this sense, the model captures markets in which the interaction is a simultaneous move game.

Second, the manufacturer has exactly two cost states and exactly two possible messages. In any equilibrium in pure strategies allowing for as many messages as states is without loss of generality, however that there are exactly two cost outcomes is an assumption made for tractability. It will be argued that for any cost the manufacturer either prefers as much search as possible or as little search as possible. With this in mind, in any cheap talk equilibrium the manufacturer may only reveal whether his cost falls into the former category or the latter. When there are only two costs, one in each category, consumers become perfectly informed and face no aggregate uncertainty. However, if either category includes more than one cost, a searching consumer continues to learn during the search process. This makes her decision problem non-stationary and thereby significantly complicates the analysis.

Lastly, consumers decide whether to accept or reject an offer but not which store to visit next. In a symmetric model this simplifying assumption is benign, and thus it is commonly used in the search literature. However, in the current environment consumers may learn whether retailers or competitors are likely to offer a better deal, and consequently may wish to direct their search. A manufacturer that signals a low cost may gain prominence in the search process, which has been shown to both directly increase his share of consumers and also potentially decrease competition (e.g. Armstrong, Vickers, and Zhou (2009), Arbatskaya (2007), Haan and Moraga-González (2011)). The random search assumption abstracts from this motive and is thus restrictive, though is justified in some situations. For instance, if a consumer is unaware of which sellers carry which products then the decision of which store to visit next is independent of which product she is likely to prefer. Alternatively, in an environment in which sellers are geographically distributed, the order of a consumer's visits may be exogenously determined by store locations.

The analysis proceeds as follows. I first solve an auxiliary game in which the manufacturer's cost is known and derive equilibrium strategies for the manufacturer, retailers, and consumers conditional on this cost. I demonstrate that when the manufacturer's cost increases, he charges a higher wholesale price and consumers are willing to accept a lower level of utility. I then use this result as well as the explicitly derived structure of the equilibrium with informed consumers to prove that the manufacturer can communicate credibly via cheap talk. That is, I derive a condition so that for each cost outcome c_L and c_H, the manufacturer prefers to induce the search behavior corresponding to the true cost rather than that associated with the other cost.

3 Equilibrium with Informed Consumers

In this section I characterize the equilibrium pricing strategies of the manufacturer and retailers and the search strategy of consumers conditional on the manufacturer truthfully reporting his costs. That is, I assume the manufacturer makes recommendations as follows:

$$\sigma(c) = \begin{cases} \sigma_L & \text{if } c = c_L \\ \sigma_H & \text{if } c = c_H \end{cases}.$$

I first solve the downstream game between retailers and consumers conditional on a manufacturer's wholesale price w, which is observed by retailers and correctly anticipated by consumers. Then, I solve for the manufacturer's optimal wholesale price.

Consumer Search and Retail Prices

First consider the behavior of consumers. Assume that search cost s is sufficiently small relative to the variance of the outside option so that searching is always optimal. That is, assume

$$s \leq \frac{M}{N+M} \int_{\underline{v}}^{\bar{v}} (v - \underline{v}) \, dG(v).$$

In deciding whether to search, a consumer evaluates the expected utility from future draws and faces two types of uncertainty. The first type includes the identity of the seller on the next draw, the realization of the outside option v_j if the seller is a competitor, and the realization of the preference shock η_i if the seller is a manufacturer's retailer. These sources of uncertainty are stationary, i.e. a consumer's belief about future realizations is independent of the consumer's previous search history. In the second type is uncertainty over the price p_i if the seller is a retailer. In the full model in which the manufacturer's cost is uncertain, previously observed retail prices can shed light on the manufacturer's cost, and in turn help refine a consumer's expectation of what prices she would face when continuing to search. However, in the present environment the consumer is informed about the manufacturer's cost and her history allows no additional inference. Hence a consumer faces a stationary decision problem and, by a standard result in Kohn and Shavell (1974), uses a stationary threshold policy $u(\sigma)$,[5] defined implicitly by

$$s = \int_{u \geq u(\sigma)} (u - u(\sigma)) \, dH(u \mid \sigma). \tag{1}$$

[5] McCall (1970) demonstrates the optimality of a stationary threshold policy when the number of draws is infinite. Kohn and Shavell (1974) extend this to a setting with finite draws, with the caveat that if the very last draw is reached, a consumer no longer uses the threshold but rather just chooses the best of the options she has seen.

The right hand side describes the option value of taking another draw, and $H(u \mid \sigma)$ is the equilibrium distribution of utilities offered by a randomly selected seller conditional on the signal σ of the manufacturer's cost realization. In an equilibrium in which the manufacturer's retailers all charge the same price $p(\sigma)$ conditional on signal σ, the distribution of utilities H can be expressed as

$$H(u \mid \sigma) = \frac{N}{N+M} F\Big(u + p(\sigma)\Big) + \frac{M}{N+M} G(u).$$

Next consider the strategy of retailers. The probability a retailer makes a sale to an arriving consumer is

$$x(p \mid \sigma) = 1 - F(u(\sigma) + p) \; + \; \Pr(\text{return}).$$

A retailer sells to a consumer either upon the consumer's arrival or if the consumer leaves but later returns. The latter occurs only if the consumer has seen all other retailers' offers and then buys from the current retailer because her utility is highest there. In particular, this requires that the maximal of all utility draws taken by a consumer is lower than her search threshold. As the number of sellers $M + N$ is increased, the probability of this event diminishes toward zero. For large values of M and N a retailer's profit and profit-maximizing price are closely approximated by ignoring the impact of returning consumers. For simplicity, in the ensuing analysis I assume that $M = N = \infty$. Thus, the probability of sale to an arriving consumer is

$$x(p \mid \sigma) = 1 - F(u(\sigma) + p). \tag{2}$$

Having observed the wholesale price w and the recommendation σ, retailer i sets price p_i and makes a sale if and only if the consumer's preference shock gives her a utility above $u(\sigma)$. Let $\theta(w, \sigma)$ denote the total measure of consumers that retailer i expects to arrive during the search process. The retailer's profit is

$$\pi(p \mid w, \sigma) = \theta(w, \sigma)(p - w)(1 - F(u(\sigma) + p))$$

and the retailer chooses a price to solve

$$p(w, \sigma) = \arg\max_{p} \; (p - w)(1 - F(u(\sigma) + p)), \tag{3}$$

which follows since $\theta(w, \sigma)$ is exogenous to the retailer and falls out from the maximization. The retailer's problem is that of a monopolist facing demand function $x(p, \sigma) = 1 - F(p + u(\sigma))$ and constant marginal cost w, and an example in which F is uniform is depicted in Figure 2.

Let $x(w, \sigma) \equiv x(p(w, \sigma) \mid \sigma)$ denote a retailer's optimally chosen probability of sale to an arriving consumer, to which I will also refer as the consumer's acceptance probability.

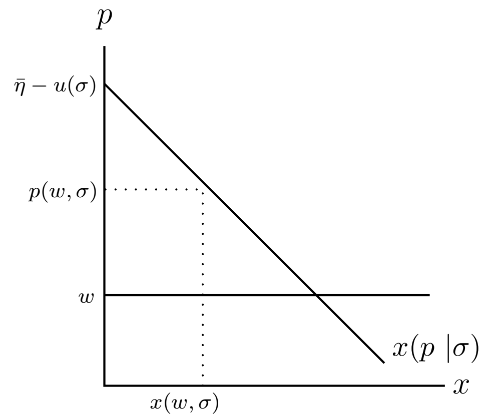

Figure 2: Retailer demand when F is uniform

Manufacturer's Wholesale Price

The quantity sold by the manufacturer is determined by the equilibrium actions of retailers and consumers in the downstream market. In the present environment the manufacturer cannot directly affect the consumers' search threshold $u(\sigma)$ but can affect retail prices through his wholesale price. The manufacturer's demand $Q(w, \sigma)$ may alternately be interpreted as the probability that an arbitrary consumer eventually purchases from a retailer and not a competitor. This interpretation allows for $Q(w, \sigma)$ to be expressed as

$$Q(w,\sigma) = \frac{1}{2}x(w,\sigma) + \left(\frac{1}{2}(1 - x(w,\sigma)) + \frac{1}{2}(G(u(\sigma)))\right)Q(w,\sigma).$$

The first term denotes the joint probability that the consumer's initial draw is at a manufacturer's retailer and that she accepts. The second term accounts for the probability that the consumer rejects her initial draw, at either type of seller, and continues to search. For her next draw the consumer faces the identical problem as for her first draw, which yields the recursive specification. Solving the above obtains

$$Q(w,\sigma) = \frac{x(w,\sigma)}{x(w,\sigma)\ +\ (1 - G(u(\sigma)))}. \tag{4}$$

The manufacturer's demand closely resembles the success function in a ratio-form (Tullock) contest, and as with the ratio-form contest the analogy of a raffle applies. A consumer using threshold $u(\sigma)$ continues to sample until reaching the first seller that meets the threshold, thus it is as if the consumer randomly samples exactly once but only from the set of qualifying sellers. The probability of sale by the manufacturer then just equals the proportion of all qualifying sellers that belong to the manufacturer.

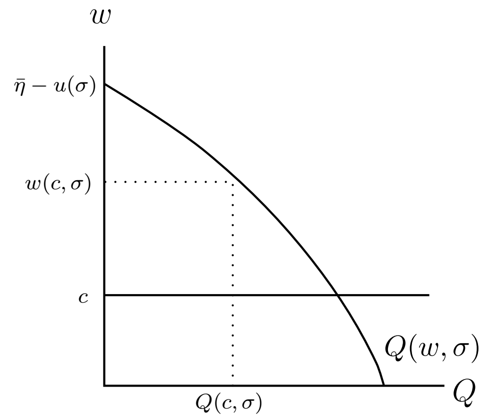

Figure 3: Manufacturer demand when F is uniform

By inspection the manufacturer's demand is concave in acceptance probability $x(w, \sigma)$, which follows solely from the search process and not from any assumption on the underlying distribution of valuations. This fact is useful for technical reasons, since the concavity of the demand function is important for the existence of a unique optimizer. It also highlights that the extent to which a manufacturer is concerned about double-marginalization depends on how his retailers compare with the competitors. In the classic setting, a one unit reduction in the sales of a retailer translates into a one unit reduction in the sales of the manufacturer. In this search environment, however, a consumer that rejects an offer from a retailer continues to search. The higher is the acceptance probability at the retailers relative to competitors, the higher the chance that a consumer induced into searching still buys from the manufacturer.[6]

Figure 3 illustrates that the manufacturer solves a monopoly problem with demand $Q(w, \sigma)$ and constant marginal cost c. His profit is

$$\begin{aligned}\Pi(w, \sigma) &= (w - c)Q(w, \sigma) \\ &= (w - c)\frac{x(w, \sigma)}{x(w, \sigma) + (1 - G(u(\sigma)))}\end{aligned}$$

and he chooses a wholesale price w which solves

$$w(c) = \arg\max_{w} \ \Pi(w, \sigma). \tag{5}$$

Proposition 1 *If $f(\eta)$ is non-decreasing then for any manufacturer cost c there exists a pure strategy equilibrium $\Big(w(c), p(w, \sigma), u(\sigma)\Big)$, which solves equations* (1), (3), *and* (5).

[6]The observation that a manufacturer is less concerned with double-marginalization when consumers search is explored in Janssen and Shelegia (2013), which demonstrates that as a consequence manufacturers set an even higher wholesale price than they would in a classic setting.

The proof relies on Schauder's fixed point theorem by demonstrating that all best responses are bounded and continuous. The continuity of the search threshold is immediate and the continuity of the retailer's profit maximizing price follows from the standard regularity assumption of log-concavity of f. The continuity of the optimal wholesale price requires more work, since the curvature of the demand faced by the manufacturer is a function of the curvature of acceptance probability $x(w, \sigma)$ with respect to w. For this result, the sufficient condition of a non-decreasing $f(\eta)$ is introduced. The full proof may be found in the Appendix.

There may be multiple equilibria in the full information setting. This owes to the complementarity of pricing decisions and search thresholds. For instance, consumers' willingness to accept an offer increases in retailer's prices and retailer's optimal prices in turn increase with consumers' willingness to accept. A similar relationship can be demonstrated for the strategies of the manufacturer and consumers. Since the central result of this paper is the existence and not uniqueness of a cheap talk equilibrium, for concreteness I focus on the equilibrium with the highest prices.

With this caveat in mind, I establish several comparative statics that are used for analyzing the cheap talk equilibrium in the main specification. First, a higher cost state is associated with a lower utility search threshold.

Lemma 1 *In the full information setting, $u(\sigma)$ is strictly decreasing in the manufacturer's marginal cost c whenever the manufacturer makes positive sales.*

Although this result appears intuitive, it does not follow immediately. While the direct effect of a higher cost is a higher wholesale price by the manufacturer, the higher wholesale price induces a lower search threshold, which can feed back to reduce the wholesale price in equilibrium. The proof, included in the appendix, shows that in any stable equilibrium the direct effect dominates the feedback effect and that the equilibrium with the highest price is stable.

Next, the consumers' search threshold increases as search cost s decreases, and approaches the maximal utility offered by competing sellers as s approaches zero.

Lemma 2 *In the full information setting, $u(\sigma)$ decreases in the consumers' search cost s and $\lim_{s\to 0} u(\sigma) = \bar{v}$.*

That the search threshold never exceeds $\bar{v}$ follows from a hold-up argument, similar to that in Janssen and Shelegia (2013) and Diamond (1971), in that a manufacturer cannot commit to provide consumers with sufficient return to searching. To see this, conjecture that the consumer uses a threshold $u > \bar{v}$. In such a situation, searching consumers accept only at retailers and never at competing sellers, which makes it costless for the manufacturer to increase his wholesale price and induce more search. The manufacturer's best response is a wholesale price at which the acceptance probability at retailers is arbitrarily close to zero, which in turn makes the expected return to searching close to zero and leads to a contradiction. The formal proof of Lemma 2 can be found in the Appendix.

4 Credibility of Manufacturer Recommendations

The preceding analysis defines search thresholds $u(\sigma_L)$ and $u(\sigma_H)$ which are consistent with consumers being informed of the manufacturer's true cost. The next step is to check whether the manufacturer can credibly provide this information via cheap talk. In such an equilibrium, regardless of the actual cost realization the manufacturer is able to induce either search threshold. It needs to be shown that the manufacturer prefers to induce $u(\sigma_L)$ when his cost is c_L and $u(\sigma_H)$ when his cost is c_H.

To build intuition for this result, consider the manufacturer's demand in equation (4) and note that when the search threshold increases, the manufacturer is both hurt by the reduced acceptance probability at his retailers and helped by the reduced acceptance probability at his competitors. Partial differentiation of the manufacturer's demand with respect to search threshold u obtains

$$\frac{\partial Q}{\partial u} = Q(w,u)(1 - Q(w,u)) \left(\frac{g(u)}{1 - G(u)} - \frac{|x_u(w,u)|}{x(w,u)} \right). \tag{6}$$

The manufacturer thus benefits from more search if and only if the acceptance probability is less elastic at his retailers than at his competitors. Which acceptance probability is more elastic in turn depends on the manufacturer's wholesale price. For example, consider a low wholesale price such that $x(w,u)$ is substantially larger than $1-G(u)$. All else equal, $x(w,u)$ is less sensitive than $1 - G(u)$ in percent terms, which in turn implies that the manufacturer sells a higher quantity when the search threshold increases. The reverse is true at a high wholesale price, in which case the manufacturer's sales fall when the search threshold increases. Figure 4 depicts how the differential effect of the search threshold can pivot the manufacturer's demand.[7]

In order for signaling to be credible, the manufacturer must prefer demand function $Q(w,\sigma_L)$ when his cost is c_L and demand function $Q(w,\sigma_H)$ when his cost is c_H. For this it is not sufficient to demonstrate that the equilibrium $w(c_L)$ is small and $w(c_H)$ is large. The argument is more nuanced because when the manufacturer misreports his cost, he can simultaneously adjust his wholesale price. For instance, when the manufacturer's cost is c_L, even if at $w(c_L)$ he prefers to report truthfully to induce more search, misreporting and inducing less search

[7]Figure 4 serves as an illustration only. For example, while the two demand curves intersect only once, and this will be shown to be the case with uniform distributions, it need not be the case in general.

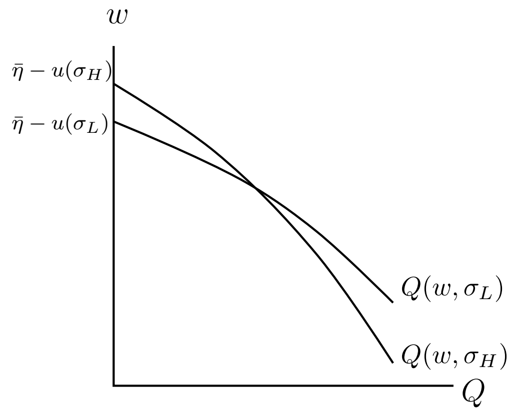

Figure 4: Effect of search when F is uniform

may be profitable if he also deviates to some higher $w > w(c_L)$. To establish the cheap talk equilibrium, it must be proven for both cost realizations that when misreporting, there exists no wholesale price that obtains a higher profit for the manufacturer than what he obtains through truth-telling.

I begin with some preliminary lemmas that simplify the analysis and provide intuition. Then, the main result that there exists a truth-telling equilibrium is demonstrated in an example in which the outside option v_j and retailer preference shocks η_i are both uniformly distributed. Finally, the result is derived in a more general setting.

Preliminaries

In order to prove the existence of an equilibrium with truthful reporting, it will be shown that at low wholesale prices the manufacturer prefers more search, and at high wholesale prices he prefers less search. Then, it will be shown that when the cost is c_L the manufacturer prefers to set a low wholesale price even when misreporting, and similarly, when the cost is c_H he sets a high wholesale price even when misreporting.

To define what constitutes a low or high wholesale price, recall from equation (6) that a key determinant of the manufacturer's preference for inducing search is the elasticity of the acceptance probability with respect to the search threshold.

Definition 1 *The magnitudes of the search elasticities of the acceptance probabilities for* *retailers and competitors are*

$$\varepsilon_x(w,u) \equiv \frac{|x_u(w(u,c),u)|}{x(w(u,c),u)}$$
$$\varepsilon_G(u) \equiv \frac{g(u)}{1-G(u)}.$$

Definition 2 *For a given search threshold $u < \bar{v}$, let $\tilde{w}(u)$ be the highest wholesale price at which search elasticities are equal for retailers and competitors. That is,*

$$\tilde{w}(u) \equiv \max\left\{\{0\} \cup \{w \mid \varepsilon_G(u) = \varepsilon_x(w,u)\}\right\}.$$

The max operator accounts for the possibility that conditional on u several wholesale prices may equalize the two elasticities, or that no wholesale price satisfies this condition. Note also that since $\tilde{w}$ is the largest wholesale price where elasticities are equal, for all $w > \tilde{w}$ the retailers' elasticity is higher since for all $w \geq \bar{\eta}$, $\varepsilon_x = \infty$. Next, the main results evaluate the effect of increased search on the manufacturer's maximized profit, and rely on the following lemma.

Lemma 3 *Let $\Pi^*(u,c)$ denote the manufacturer's maximized profit and $w^*(u,c)$ the profit-maximizing wholesale price when consumers use search threshold u and the marginal cost is c, and let $u_2 > u_1$. Then*

$$\Pi^*(u_2,c) - \Pi^*(u_1,c) = \int_{u_1}^{u_2} Q(w^*(u,c),u)\left(\frac{\varepsilon_G(u)}{\varepsilon_x(w^*(u,c),u)} - 1\right)\, du.$$

The proof is included in the Appendix and makes use of the envelope theorem to circumvent solving directly for the profit maximizing wholesale price. This is particularly useful in light of the fact that to check whether misreporting is profitable, one must evaluate profit at the optimal deviation wholesale price, for which no closed form solutions exists. This formulation implies the following corollary, which provides an easily stated sufficient condition for the manufacturer's preference for search based solely on search elasticities.

Corollary 2 *If $\varepsilon_G(u) > (<)\ \varepsilon_x(w^*(u,c),u)$ for all $u \in [u_1,u_2]$ then $\Pi^*(u_2,c) - \Pi^*(u_1,c) > (<)\ 0$.*

It will be sufficient for the cheap talk result to show that for any search threshold $u \in [u_1,u_2]$, the manufacturer would induce low elasticities ε_x when cost is c_L and high elasticities ε_x when cost is c_H. To ascertain whether the manufacturer prefers to induce a higher or lower elasticity, it is useful to determine the manufacturer's marginal profit at the wholesale price

$\tilde{w}$ at which the two elasticities are equal. Differentiating the manufacturer's profit function and evaluating at $\tilde{w}(u)$ simplifies to

$$\frac{\partial \Pi(\tilde{w}(u), u)}{\partial w} = Q(\tilde{w}(u), u)\left(1 - (\tilde{w}(u) - c)\left(\frac{g(u)}{g(u) + x_u(\tilde{w}(u), u)}\right)\varepsilon_G(u)\right). \tag{7}$$

Since assumptions on the distribution $F(\eta)$ guarantee that the manufacturer's profit is concave, the sign of marginal profit determines whether for a given u the manufacturer sets a wholesale price above or below $\tilde{w}(u)$, hence whether he optimally induces a higher or lower search elasticity of acceptance for retailers than for consumers.

The strategy to prove the existence of a cheap talk equilibrium is to use equation (7) to show that for any $u \in [u(\sigma_H), u(\sigma_L)]$ the manufacturer sets a low wholesale price below $\tilde{w}$ when his cost is c_L and a high wholesale price above $\tilde{w}$ when his cost is c_H. Then Corollary 2 is used to show that truth-telling is a best response in both states. The result is shown when utility shocks are uniformly distributed, and then in a more general setting with some additional regularity conditions.

Uniform Example

Let the distribution of the outside option G be uniform on the interval $[\underline{v}, \bar{v}]$ and the distribution of the preference shocks at the manufacturer's retailers F be uniform on the interval $[\underline{\eta}, \bar{\eta}]$. Let $\bar{\eta} - c_H < \bar{v} < \bar{\eta} - c_L$, which ensures that the maximum utility provided by competitors is lower than the maximum that the manufacturer could provide in the low cost state, and higher than what he could provide in the high cost state.

First consider the retailers' problem. With uniform utility shocks, each retailer is a monopolist facing linear per-visitor demand

$$x(p \mid \sigma) = 1 - \frac{1}{\bar{\eta} - \underline{\eta}}(\bar{\eta} - u(\sigma) - p).$$

Solving the profit maximization from equation (3) yields acceptance probability

$$x(w, \sigma) = \frac{1}{2(\bar{\eta} - \underline{\eta})}(\bar{\eta} - u(\sigma) - w). \tag{8}$$

Next, the expression for the acceptance probability in equation (8) is used to solve for the wholesale price $\tilde{w}(u)$ which equates search elasticities across retailers and competitors.

$$\varepsilon_x(\tilde{w}(u), u) = \varepsilon_G(u)$$
$$\Leftrightarrow$$
$$\tilde{w}(u) = \bar{\eta} - \bar{v} \tag{9}$$

Note that in the uniform specification $\tilde{w}(u)$ is constant with respect to u, though in general this is not the case. Having obtained expressions for the acceptance probability and thus search elasticity as well as $\tilde{w}(u)$, plugging into the manufacturer's marginal profit evaluated at $\tilde{w}(u)$ using expression (7) obtains

$$\frac{\partial \Pi(\tilde{w}(u), u)}{\partial w} = Q(\tilde{w}(u), u)\left(1 - ((\bar{\eta} - c) - \bar{v})\left(\frac{2(\bar{\eta} - \underline{\eta})}{2(\bar{\eta} - \underline{\eta}) \quad + \quad (\bar{v} - \underline{v})}\right)\frac{1}{\bar{v} - u}\right). \tag{10}$$

The sign of the marginal profit is determined by the sign of the term in the outer parentheses. The fact that u is the only endogenous object in this term simplifies the analysis, and leads to the following proposition.

Proposition 3 *With uniform utility shocks, for a sufficiently small search cost s, there exists a cheap-talk equilibrium in which the manufacturer truthfully reveals his costs.*

Proof The aim is to evaluate the sign of equation (10), that is

$$sign(\Pi_w(\tilde{w}(u), u)) = sign\left(1 - ((\bar{\eta} - c) - \bar{v})\left(\frac{2(\bar{\eta} - \underline{\eta})}{2(\bar{\eta} - \underline{\eta}) \quad + \quad (\bar{v} - \underline{v})}\right)\frac{1}{\bar{v} - u}\right)$$

Recalling that $\bar{\eta} - c_H < \bar{v}$ by assumption, observe first that when $c = c_H$ the marginal profit is positive at $\tilde{w}(u)$ for any $u < \bar{v}$. In other words, when the manufacturer's cost is high he optimally sets a high enough wholesale price to induce a larger search elasticity for retailers than for competitors. This in turn implies by Corollary 2 that the manufacturer strictly prefers a lower search threshold, which is achieved by reporting σ_H.

When $c = c_L$ the sign of marginal profit at $\tilde{w}(u)$ is ambiguous, however it will now be argued that the sign is negative when s is sufficiently small. First, fix c_L and c_H and note that by Lemma 2 that as s is reduced toward zero, both thresholds $u(\sigma_H)$ and $u(\sigma_L)$ approach $\bar{v}$. Consequently, for small enough s, for any $u \in [u(\sigma_H), u(\sigma_H)]$, the difference $\bar{v} - u$ is arbitrarily small, resulting in a negative marginal profit in equation (10). Thus, by Corollary 2, for small values of s the manufacturer strictly prefers a higher search threshold and reports σ_L.

It is thus a best response for the manufacturer to report truthfully in both states. For completeness, the consumers' strategy must be revisited in this Bayesian game. A consumer that has seen k retailers and l competitors has a history

$$h = \{\sigma, \{(p_1, \eta_1), ..., (p_k, \eta_k)\} \cup \{v_1, ..., v_l\}\} = \{\sigma, v\},$$

and her strategy is a mapping from histories to a set of actions which includes accepting an available offer, searching, or exiting. With each history is associated a posterior $\mu(h)$ about the probability that the manufacturer has a high cost, and this posterior must be formed using Bayes rule whenever possible. Let the consumers' posterior be

$$\mu(h) = \begin{cases} 1 & \text{if } \sigma_H \in h \\ 0 & \text{if } \sigma_L \in h \end{cases}.$$

These beliefs are consistent with equilibrium strategies, but impose a restriction off the equilibrium path. Namely, if a consumer ever observes a retail price that is not consistent with the recommended price, the consumer places all the weight on the price recommendation.[8]

Thus, there exists a sequential equilibrium in the Bayesian game in which the manufacturer uses signals $\sigma(c_H) = \sigma_H$ and $\sigma(c_L) = \sigma_L$, consumers have beliefs given by $\mu(h)$ above, and strategies are described by those that solve the full information game in Proposition 1. ■

General Result

Assume the support of distribution F is bounded above by $\bar{\eta}$ and the support of distribution G is bounded above by $\bar{v}$. This ensures that as the search cost is reduced toward zero, consumers' do not demand an infinite utility. In place of the uniform distributions as above, assume that densities g and f and marginal density f' are all bounded. Assume again that $\bar{\eta} - c_H < \bar{v} < \bar{\eta} - c_L$, so that the maximal gains from trade between the manufacturer and consumers is higher than the maximal outside option utility when his cost is low, and the reverse when his cost is high.

To prove that the manufacturer prefers to report truthfully, I use the same technique as above by examining the sign of the marginal profit in equation (7). While in the uniform setting, the sign of this equation reduced in equation (10) to a function of only parameters and the search threshold u, presently other equilibrium objects such as wholesale price $\tilde{w}(u)$ and the marginal acceptance probability $x_u(\tilde{w}(u), u)$ remain, requiring additional analysis.

Lemma 4 *As the search cost s approaches zero, $\tilde{w}(u)$ approaches $\bar{\eta} - \bar{v}$.*

Lemma 2 already established that $\lim_{s \to 0} u = \bar{v}$, thus as s approaches zero the search elasticity ε_G at competitors approaches infinity. By definition, at the wholesale price $\tilde{w}$ the

[8] These beliefs are sufficient but not necessary to support the cheap talk equilibrium. While these particular beliefs were chosen for simplicity, they may also be justified by considering a related model in which retailers have private idiosyncratic costs, which in turn would induce a distribution of retail prices for each price recommendation. In such a setting the set of off-the path prices is reduced and the assumption becomes less restrictive. An additional defense is the main result itself, namely that manufacturers' and consumers' interests about consumers' search decisions are aligned, while retailers all would prefer less search. By a logic similar to the intuitive criterion, consumers ought to trust the manufacturer more than retailers.

search elasticity for retailers must also approach infinity, which implies that the retailers' acceptance probability must approach zero and thus $\tilde{w}$ must also approach $\bar{v}$. The proof relies on the assumptions of bounded densities and is included in the Appendix. Note that the uniform $\tilde{w}$ from equation (9) is a special case.

Proposition 4 *For a sufficiently small search cost s, there exists a cheap-talk equilibrium in which the manufacturer truthfully reveals his costs.*

Proof The aim is to evaluate the sign of equation (7), that is

$$sign(\Pi_w(\tilde{w}(u), u)) = sign\left(1 - (\tilde{w}(u) - c)\left(\frac{g(u)}{g(u) + x_u(\tilde{w}(u), u)}\right)\varepsilon_G(u)\right)$$

First consider when $c = c_H$. By Lemma 4, $\lim_{s\to 0} \tilde{w} = \bar{\eta} - \bar{v} < c_H$ implies that for small s the marginal profit is strictly positive. Hence, for all $u \in [u(\sigma_H, \sigma_L]$ the manufacturer sets a wholesale price $w > \tilde{w}(u)$, inducing a higher elasticity for his retailers than for competitors. He prefers less search and thus reporting σ_H is a best response.

Next consider when $c = c_L$. Since $\lim_{s\to 0} \tilde{w} = \bar{\eta} - \bar{v} > c_L$, the markup $\tilde{w} - c_L$ is bounded above zero. The term $\frac{g(u)}{g(u)+x_u(\tilde{w}(u),u)}$ is bounded since $g(\cdot)$ is bounded by assumption and x_u is shown to be bounded in the proof of Lemma 4. A bounded $g(u)$ also implies that $\lim_{s\to 0} \varepsilon_G(u) = \lim_{s\to 0} \frac{g(u)}{1-G(u)} = \infty$. Thus, for small s, $sign(\Pi_w(\tilde{w}(u), u)) < 0$. This implies that for all $u \in [u(\sigma_H, \sigma_L]$ the manufacturer sets a wholesale price $w < \tilde{w}(u)$, inducing a lower elasticity for his retailers than for competitors. He prefers more search and thus reporting σ_L is a best response.

Therefore, it is a best response for the manufacturer to report truthfully in both states. For completeness, consumers' off the path beliefs are the $\mu(h)$ specified in Proposition 3. ■

The results in Propositions 3 and 4 demonstrate that the manufacturer can use cheap talk to inform consumers. However, the same equilibrium exists even if the manufacturer cannot communicate at all. Specifically, in the equilibrium with communication there is a single retail price associated with the low cost state and a single retail price associated with the high cost state. A consumer can perfectly infer the cost realization after observing one retail price, thus obviating manufacturer recommendations. This observation however is more technical than it is qualitative, as the existence of an equivalent equilibrium without manufacturer signaling is not robust. For instance, consider a modification to the model in which retailers have idiosyncratic cost shocks in addition to the wholesale price. In equilibrium, all retailers face the same demand function and due to cost differences choose different monopoly prices. A consumer must now infer from a retail price whether it reflects the retailer's or the manufacturer's cost shock, and in general will not perfectly learn the latter.

5 Discussion

Having described the equilibrium in which manufacturers use recommendations to inform consumers, I draw the comparison to an environment in which recommendations are not made, for instance if they are banned by an antitrust authority. Along similar lines, since MSRPs have been viewed as a means for manufacturers to exert influence over retail prices, I contrast their effect to the effect of maximum resale price maintenance.

Comparison to No Communication

The question of what would ensue if the manufacturer does not make recommendations is important for two reasons. First, as with any cheap talk environment, there always exists a babbling equilibrium in which consumers ignore manufacturer recommendations and recommendations are uncorrelated with the realization of manufacturer costs (for instance, the signal "no recommendation" is always used). Given that in most markets it is optional[9] and often costly for the manufacturer to print price recommendations, it helps to argue that the informative equilibrium is more profitable for the manufacturer than the babbling equilibrium. Second, being able to compare the informative equilibrium with an equilibrium without MSRPs is important from a policy perspective in determining whether manufacturer recommendations should be regulated.

An explicit comparison of the two environments – with and without recommendations – is made difficult by the complexity of solving the latter. This stems from the fact that when consumers are uncertain about the distribution of utilities from which they sample, their optimal search strategy is no longer easily described by a single threshold. Instead, the strategy is non-stationary, so that a consumer's willingness to accept a price is a function of her search history. In order to find a profit maximizing price, a retailer must form a posterior of a visiting consumer's history, and furthermore must anticipate the consumer's future draws to compute the probability of her return. Solving the consumers' and retailers' problems simultaneously is quite difficult and remains an open question in the search literature.[10]

Despite this difficulty in characterizing an explicit solution, the construction of the informative equilibrium in Section 4 sheds some light on the factors that would affect manufacturers and consumers if MSRPs were removed. First, conjecture that when consumers receive no

[9] An important counterexample in which an MSRP is mandatory is the Monroney sticker affixed to all new cars in the United States, the result of the Automobile Information Disclosure Act of 1958. Even in this case however, it is not clear that the MSRP must be informative. A manufacturer is in principle allowed to set the same price recommendation for two vehicles he expects to sell for different prices.

[10] To circumvent the issue of non-stationarity in models of search with aggregate uncertainty, two approaches have been pursued in the literature. Papers have used non-sequential search (Yang and Ye (2008), Tappata (2009)), which imposes stationarity by construction, or used sequential search with exactly two firms (Benabou and Gertner (1993), which reduces the set of relevant histories.

signal, they follow a search strategy in which they are on average less selective than if they were fully informed in the low cost state and more selective than if they were fully informed in the high cost state. Since the manufacturer tends to prefer more search when his cost is c_L and less search when his cost is c_H, in either state moving toward an intermediate amount of search reduces his payoff. It is thus likely that the manufacturer is better off when he is allowed to make price recommendations.

For consumers, the direct effect of receiving information through manufacturer recommendations is beneficial since it fosters better search decisions. However, there is also an equilibrium effect in that the set of options available to consumers changes depending on whether MSRPs are present. In general, consumers prefer to commit to more search since this lowers prices. If on average consumers search more when uninformed than when informed, then from this they derive an additional benefit from MSRPs beyond better decision-making. However, if uncertainty makes consumers search less on average then this could lead to higher prices and potentially dominate the benefit of better information. In a related setting, Benabou and Gertner (1993) demonstrates that a reduction in uncertainty can either help or harm searching consumers in equilibrium, depending on market parameters.

Comparison to Resale Price Maintenance

In practice retail prices rarely exceed recommendations and consequently MSRPs are often viewed as price ceilings. With explicit price ceilings, a retailer faces a penalty for setting a price above the ceiling that is typically enforced by the manufacturer. Similarly, MSRPs have been conjectured to provide a punishment for retailers but through consumers who refuse to pay above the recommended price, or more generally have a kink in their demand at this price (Thaler (1985), Puppe and Rosenkranz (2011)).

Such theories predict that the presence of MSRPs reduces demand only at prices that exceed the recommendation. In contrast, the present theory of MSRPs as information for consumers predicts that demand is potentially affected at all prices, including those below MSRP. Specifically, in the model the per-person demand is $x(p, u) = 1 - F(p + u)$, thus an increase in search threshold u corresponds to an inward shift of demand for every price.

Little direct evidence exists to determine which of the two predictions is borne out empirically,[11] however the US automobile market provides an anecdotal example. The vast majority of cars sell not just below MSRP, but strictly below, so that in most transactions the MSRP does not bind. In principle, it is possible that consumers use some fraction of MSRPs as a reference point, so for instance the true price ceiling is at 95% of the recommended

[11] In a recent paper, De los Santos et. al. (2013) examines the effect of an MSRP ban and later reinstatement on prices and finds that while the presence of MSRPs reduces prices, there is no evidence of recommendations acting as binding price ceilings.

price. However, it is also commonly known that more popular models sell closer to MSRP, and very popular models may even sell above MSRP. This necessitates that consumers have market-dependent reference points, which, essentially is equivalent to this paper's model in which consumers rationally condition their search thresholds based on their assessment of the distribution of prices for a particular product.

As a conduit of information MSRPs provide the manufacturer with an instrument that is both similar to and different from resale price maintenance. On one hand, MSRPs affect consumers' search thresholds and thereby indirectly affect the prices set by retailers. In this sense, the sending of signals σ_L and σ_H may be interpreted as a way to reduce or increase retail prices, as can be achieved by resale price maintenance. On the other hand, price recommendations affect the behavior of consumers when they visit competing firms. By sending signal σ_L for example, the manufacturer reduces demand at this competitors at every price. Through this channel, MSRPs can be used to increase or soften competition.

6 Conclusion

Despite the prevalence of manufacturer suggested retail prices, why manufacturers make recommendations and what effect they have on market outcomes such as prices and sales is still not well understood. I present a model in which a manufacturer uses recommendations to affect consumers' search decisions. Consumers do not know the distribution of retail prices due to uncertain manufacturer costs and rely on the recommendation to determine whether to purchase at a given price or to continue shopping. By using a suggested price to induce more search the manufacturer pressures retailers to reduce double marginalization but also loses consumers that search and purchase from a competitor. The manufacturer profits from more search when he offers consumers more value than competitors; otherwise he prefers less search. Consumers benefit from more search when higher utilities are available, thus their interests are aligned with the manufacturer's. I demonstrate that manufacturer recommendations can be used as cheap-talk signals to credibly inform consumers of the manufacturer's costs. Suggested prices help manufacturers and consumers coordinate and thereby can improve the welfare of both.

It is likely that the manufacturer's incentive for truthful reporting applies beyond the current setting of uncertainty over costs. For example, suppose the manufacturer's cost is fixed and instead there is uncertainty about product quality. In terms of the current model, suppose that the distribution of utility shocks has an uncertain mean α, with low realization α_L and high realization α_H. The manufacturer learns α and can communicate it with a cheap message. If retailers do not fully collect the surplus from higher quality, consumers ought to use a higher utility threshold when $\alpha = \alpha_H$ and a lower threshold when $\alpha = \alpha_L$. Meanwhile, the manufacturer prefers more search when $\alpha = \alpha_H$ because he offers a higher distribution of utilities than his competitors, and similarly prefers less search when $\alpha = \alpha_L$.

In principle it is difficult for a seller to convey the value of his product to a buyer using cheap talk. The intuition is simple – if a consumer's willingness to pay for a product can be increased by a cheap message, this message is used regardless of the product's value (e.g. Milgrom and Roberts (1986)). I demonstrate that in a search environment this intuition no longer holds; instead the manufacturer can affect consumers' willingness to pay by informing them of the set of other options available through search. In particular, the manufacturer can sometimes benefit by informing consumers that the product he supplies retailers is not of high quality on average, thereby reducing the perceived return to searching and making them more willing to accept the current offer. Since the manufacturer does not always benefit from overstating the value of his product, he is able to communicate credibly.

7 Appendix

Proof of Proposition 1: Existence of equilibrium with informed consumers

The proof relies on Schauder's fixed point theorem. The following series of lemmas establishes that the best responses of the manufacturer, retailers, and consumers come from a compact space, and then that each best response is continuous in the actions of the other players.

Lemma 5 *Let the smallest solutions of equations* (1)*,* (3)*, and* (5) *define* $\mathcal{M} : \mathbb{R}^3_+ \to \mathbb{R}^3_+$*. If* $(u, p, w) \in [s, \bar{\eta}] \times [0, \bar{\eta}] \times [c, \bar{\eta}]$ *then* $\mathcal{M}(u, p, w) \in [s, \bar{\eta}] \times [0, \bar{\eta}] \times [c, \bar{\eta}]$.

Proof of Lemma: First consider the consumers' threshold in equation (1). By assumption the returns to searching from only distribution $G(\cdot)$ outweigh the search cost, hence $u(\sigma) \geq s$. In the other extreme, for any positive retail price $u(\sigma) = \bar{\eta}$ strictly dominates any larger threshold.

Next, the profit maximizing retail price from (3) must be no smaller than the wholesale price, and any price above $\bar{\eta}$ guarantees zero profit, which can always be achieved with a price $p = \bar{\eta}$.

Lastly, the profit maximizing wholesale price from (5) must be no lower than the marginal cost. On the other hand, if it is optimal to set a wholesale price above $\bar{\eta}$ and thus make no sales, it is also optimal to set $w = \bar{\eta}$. ■

The next step is to demonstrate that $\mathcal{M}$ is continuous. To simplify notation, since in equilibrium $u(\sigma)$ is a one-to-one mapping all objects that condition on signal σ will be written to condition on u.

Lemma 6 *The optimal wholesale price w from equation* (5) *is unique and continuous in p and u whenever f is non-decreasing.*

Proof of Lemma: The manufacturer's profit function in (5) is continuous in p and u, hence it will be sufficient to demonstrate that profit is concave in w. This will be done by first showing that the manufacturer's demand is concave in the acceptance probability (Claim 1) and then that the acceptance probability itself is concave in w (Claim 2).

Claim 1 *If f is non-decreasing then* $\frac{\partial^2 x(w,u)}{\partial w^2} < 0$.

Proof of Claim: Recall that the probability of sale by a retailer is given by $x(p, u) = 1 - F(u+p)$ and define the inverse demand as $A(x) \equiv F^{-1}(1-x) - u$. Rewrite the retailer's problem.

$$\pi(x) = x(A(x) - w)$$
$$\pi'(x) = 0 = A + xA_x - w$$

The first order condition implicitly defines optimal acceptance probability $x(w, u)$. Differentiating with respect to w yields

$$x_w(2A_x - xA) - 1 = 0.$$

Differentiating with respect to w one more time yields

$$x_{ww} + (x_w)^3(3A_{xx} + xA_{xxx}) = 0.$$

Note that $x_w < 0$ because the retailer's profit has increasing differences in p and w. Then, by inspection $x_{ww} < 0$ whenever $3A_{xx} + xA_{xxx} < 0$. Using the definition of $A(x)$, this condition translates to

$$-\frac{f'}{f^3} < \frac{x}{3}\left(3\frac{|f'|}{f} - \frac{f''}{|f'|}\right).$$

By the log-concavity of f the right hand side is positive, thus whenever $f' \geq 0$ the inequality holds.

Claim 2 *If* $\frac{\partial^2 x}{\partial w^2}(w, u) < 0$ *then there exists a unique profit maximizing w for the manufacturer, which is continuous function $w^*(u)$) which is the unique solution to equation* (5).

Proof of Claim: For notation, let $Q_w \equiv \frac{dQ(w,u)}{\partial w}$ and $x_w \equiv \frac{dx(w,u)}{dw}$. First, I show that the manufacturer's demand is concave in w.

$$Q(w, u) = \frac{x(w, u)}{x(w, u) \;+\; (1 - G(u))}$$
$$Q_w = x_w \frac{1 - G(u)}{(x(w, u) \;+\; (1 - G(u)))^2}$$
$$Q_{ww} = Q_w\left(\frac{-2x_w}{x(w, u) \;+\; (1 - G(u))} \;+\; \frac{x_{ww}}{x_w}\right)$$

Since the retailer's payoff has increasing differences in w and p, $x_w < 0$ and consequently $Q_w < 0$. By inspection the sign of the expression above is negative whenever $x_{ww} < 0$. Thus when this condition is satisfied, demand is concave which ensures a unique continuous maximizer. ∎

Lemma 7 *The mapping $\mathcal{M}$ is continuous whenever f is non-decreasing.*

Proof of Lemma: It has been demonstrated that w is continuous. Next, the retailer's profit maximizing price from equation (3) is unique and continuous by the assumption of log-concavity of f. Finally, by inspection of equation (1) the consumer's best response threshold u is continuous in p, and by extension in w which has a continuous indirect effect through p. ∎

By Lemmas 5 and 7, using Schauder's theorem the mapping

$$\mathcal{M} : [s, \bar{\eta}] \times [0, \bar{\eta}] \times [c, \bar{\eta}] \;\rightarrow\; [s, \bar{\eta}] \times [0, \bar{\eta}] \times [c, \bar{\eta}]$$

has a fixed point, thus there exists a pure strategy equilibrium, which concludes the proof of Proposition 1. ∎

Proof of Lemma 1: Search threshold u is decreasing in manufacturer cost c

The intuition behind the result is that a higher cost induces the manufacturer to charge a higher wholesale price, which in turn makes consumers more pessimistic about searching. The concern is that the feedback effect of a reduced consumer search threshold on the optimal wholesale price may overpower the direct effect of the cost increase. It will be demonstrated below that the direct effect dominates.

Let $w(u, c)$ be the manufacturer's best response function, defined in equation (5) as

$$w(u, c) = \arg\max_{w} \; (w - c)\frac{x(w, u)}{x(w, u) + (1 - G(u))}.$$

Whenever $u > c$ the solution has been shown to be unique, however if $u \leq c$ then any w at which no sales are made is a best response. In these cases let $w(u, c) = c$.

Let $u(w)$ be the consumer's best response function implicitly defined in equation (1) as

$$s = \frac{1}{2}\int_{v \geq u(w)} (v - u(w))dG(v) \; + \; \frac{1}{2}\int_{\eta \geq u(w) + p(w, u(w))} (\eta - p(w, u(w)) - u(w))dF(\eta),$$

and let $\varphi(u)$ be the consumers' inverse best response function. An equilibrium in the full information setting occurs whenever $w(u,c) = \varphi(u)$, and totally differentiating the equilibrium condition with respect to c obtains:

$$\frac{d}{dc}(w(u(c),c)) = \frac{d}{dc}\varphi(u(c))$$
$$\frac{\partial w}{\partial u}\frac{du}{dc} + \frac{\partial w}{\partial c} = \frac{\partial \varphi}{\partial u}\frac{du}{c}$$
$$\frac{du}{dc} = -\frac{\frac{\partial w}{\partial c}}{\frac{\partial w}{\partial u} - \frac{\partial \varphi}{\partial u}}$$

First note that $\frac{\partial w}{\partial c} > 0$ since the manufacturer's payoff has increasing differences in w and c, which in turn follows because the retailers' optimal acceptance probability $x(w,u)$ is decreasing in w.

Therefore, in order for $u(c)$ to be decreasing, it must be that at an equilibrium $\frac{\partial w}{\partial u} - \frac{\partial \varphi}{\partial u} < 0$. To show this, since both functions are continuous it is sufficient to demonstrate that there exist $u_1 < u_2$ such that $w(u_1,c) < \varphi(u_1)$ and $w(u_2,c) \geq \varphi(u_2)$.

First let u_1 be the solution to $s = \frac{1}{2}\int_{u \geq u_1}(u - u_1)dG(u)$, thus the threshold used when a consumer expects to accept offers only from competitors. For this to be optimal, consumers must expect a high enough wholesale price such that no retailer can ever meet the search threshold, hence $\varphi(u_1) \geq \bar{\eta}$. On the other hand, the manufacturer's optimal price $w(u_1,c) \in (c,u_1)$, that is he should make sales with positive probability whenever he is able to do so. By construction $u_1 < \bar{\eta}$, thus $w(u_1,c) < \varphi(u_1)$.

Next, choose u_2 such that $\varphi(u_2) = c$, that is the threshold consumers would use when expecting the manufacturer to set $w = c$. If $u_2 < c$ then the manufacturer is unable to make positive profit and one best response is to set $w = c$. If $u_2 \geq c$ then the manufacturer's best response is unique and $w \geq c$. Thus, $w(u_2,c) \geq \varphi(u_2)$.

Note that this type of equilibrium, in which $\frac{\partial w}{\partial u} - \frac{\partial \varphi}{\partial u} < 0$ is stable with respect to best response dynamics, while any equilibrium in which the reverse is true is unstable. Note also that when $c \geq u$, the manufacturer makes no sales and $\frac{\partial w}{\partial c} = 0$, implying that $\frac{du}{dc} = 0$. However, for any $c < u$, $\frac{\partial w}{\partial c} > 0$ and consequently $\frac{du}{dc} < 0$. ■

Proof of Lemma 2: Search threshold u is increasing in search cost s and $\lim_{s\to 0} u(\sigma) = \bar{v}$

The direct effect of an increase in s is a lower search threshold, causing an inward shift of the retailers' demand function. That u decreases in s in equilibrium follows from the same technique as the proof of Lemma 1.

To demonstrate that the threshold approaches $\bar{v}$ as s approaches zero, fix a $\delta > 0$ and let

$$s = \frac{1}{2}\int_{v \geq \bar{v}-\delta} (v - (\bar{v} - \delta))dG(v),$$

where the right hand side is the value from searching when using threshold $\bar{v} - \delta$ and only accepting offers from competitors. Fix $s \in (0, s)$. Then it follows that

$$\int_{u \geq u(s')} (u - u(s'))dH(u) = s' < s \leq \int_{u \geq \bar{v}-\delta} (u - (\bar{v} - \delta)dH(u).$$

The final inequality follows from the definition of the search threshold in equation (1), using the fact the return to searching when accepting only at competitors is weakly smaller than when accepting from all sellers. Comparing the first and last terms, it follows that $u(s') > \bar{v} - \delta$. ∎

Proof of Lemma 3: Maximized profit with respect to the search threshold

The lemma claims that

$$\Pi^*(u_2, c) - \Pi^*(u_1, c) = \int_{u_1}^{u_2} Q(w^*(u, c), u)\left(\frac{\varepsilon_G(u)}{\varepsilon_x(w^*(u, c), u)} - 1\right) du.$$

First, the rate of change of maximized profit with respect to search threshold u is given by

$$\frac{d\Pi^*(u)}{du} = \frac{d\Pi(w^*(u, c), u)}{du} = \frac{\partial\Pi(w^*, u)}{\partial w^*}\frac{\partial w^*(u, c)}{\partial u} + \frac{\partial\Pi(w^*, u)}{\partial u}$$

By the envelope theorem, the first term in the sum is zero, thus

$$\begin{aligned}\frac{d\Pi^*(u)}{du} &= \frac{\partial\Pi(w^*, u)}{\partial u}\\ &= (w^* - c)\frac{\partial Q(w^*(u, c), u)}{\partial u}\\ &= (w^* - c)Q(w^*, u)(1 - Q(w^*, u))(\varepsilon_G(u) - \varepsilon_x(w^*, u)) \qquad (11)\end{aligned}$$

Above, the final line follows from equation (6). The expression can be simplified further by using the manufacturer's first order condition with respect to the wholesale price.

$$\begin{aligned}\frac{d\Pi(w^*,u)}{dw} &= 0 = Q(w^*,u) + (w-c)\frac{\partial Q(w^*,u)}{\partial w} \\ &= Q(w^*,u)\left(1 + (w^*-c)(1-Q(w^*,u))\frac{x_w(w^*,u)}{x(w^*,u)}\right) \\ (w^*-c)(1-Q(w^*,u)) &= -\frac{x(w^*,u)}{x_w(w^*,u)} \\ &= \varepsilon_x(w^*,u)\end{aligned}$$

The last equality follows from the fact that $x_w(w^*,u) = x_u(w^*,u)$, which may be derived from the retailer's first order condition.[12] Plugging back into (11) obtains

$$\frac{d\Pi^*(u)}{du} = Q(w^*,u)\left(\frac{\varepsilon_G(u)}{\varepsilon_x(w^*,u)} - 1\right).$$

Consequently,

$$\begin{aligned}\Pi^*(u_2,c) - \Pi^*(u_1,c) &= \int_{u_1}^{u_2} \Pi_u^*(u)\ du \\ &= \int_{u_1}^{u_2} Q(w^*(u,c),u)\left(\frac{\varepsilon_G(u)}{\varepsilon_x(w^*(u,c),u)} - 1\right)\ du.\end{aligned}$$

■

Proof of Lemma 4: $\lim_{s\to 0}\tilde{w}(u) = \bar{\eta} - \bar{v}$

Lemma 2 already established that $\lim_{s\to 0} u = \bar{v}$. Since g is bounded, this implies that $\lim_{s\to 0}\varepsilon_G(u) = \lim_{s\to 0}\frac{g(u)}{1-G(u)} = \infty$. By definition of $\tilde{w}(u)$, this in turn implies that $\lim_{s\to 0}\varepsilon_F(\tilde{w}(u),u) = \lim_{s\to 0}\frac{x_u(\tilde{w}(u),u)}{x(\tilde{w}(u),u)} = \infty$.

Next, derive an expression for $\frac{x_u(\tilde{w}(u),u)}{x(\tilde{w}(u),u)}$ in terms of the distribution of preference shocks F by starting with the retailer's first order condition from equation (3).

$$\frac{d\pi(p,w)}{dp} = 0 = 1 - F(p(u,w)+u) - (p(u,w)-w)\cdot f(p(u,w)+u)$$

[12] As in Claim 1, a retailer's inverse demand is $A(x) \equiv F^{-1}(1-x) - u$ and his profit as a function of the acceptance probability is

$$\begin{aligned}\pi(x) &= x(A(x)-w) \\ &= x(F^{-1}(1-x) - (u+w))\end{aligned}$$

Since $\frac{d\pi}{du} = \frac{d\pi}{dw}$, the optimizer $x(w,u)$ is affected equally by both parameters.

Suppressing the arguments of the profit-maximizing price, differentiation with respect to u obtains

$$\frac{dp}{du} = \frac{f(p+u)}{f(p+u) + 2(p-w)f'(p+u)} - 1$$

This, in turn, translates into the rate of change of the profit maximizing quantity

$$\frac{dx}{du} = \frac{f(p+u)}{f(p+u) + 2(p-w)f'(p+u)}.$$

The retailers' search elasticity of acceptance is thus

$$\varepsilon_x(\tilde{w}(u), u) = \frac{x_u(\tilde{w}(u), u)}{x(\tilde{w}(u), u)} = \frac{1}{1 - F(p+u)} \; \frac{f(p+u)}{f(p+u) + 2(p-\tilde{w})f'(p+u)}$$

Note that since f and f' are both bounded and $f(\cdot)$ is non-decreasing, the second fraction is bounded both from above and away from zero. Since ε_x must approach infinity, $1-F(p+u)$ must approach zero, which means $p+u$ must approach $\bar{\eta}$. This can only occur if $\tilde{w}$ approaches $\bar{\eta}-\bar{v}$, so that retailers' marginal cost approaches the consumers' maximal willingness to pay.
■

References

Anderson, S. P., and R. Renault (1999): "Pricing, product diversity, and search costs: a Bertrand-Chamberlin-Diamond model," *The RAND Journal of Economics*, 30(4), pp. 719–735.

Arbatskaya, M. (2007): "Ordered search," *The RAND Journal of Economics*, 38(1), 119–126.

Armstrong, M., and Y. Chen (2012): "Discount pricing," *Working Paper*.

Armstrong, M., J. Vickers, and J. Zhou (2009): "Prominence and consumer search," *The RAND Journal of Economics*, 40(2), 209–233.

Benabou, R., and R. Gertner (1993): "Search with learning from prices: does increased inflationary uncertainty lead to higher markups," *The Review of Economic Studies*, 60(1), pp. 69–93.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile prices in market equilibrium," *Econometrica*, 63(4), pp. 841–890.

Buehler, S., and D. L. Gärtner (2013): "Making sense of nonbinding retail-price recommendations," *The American Economic Review*, 103(1), pp. 335–359.

Crawford, V. P., and J. Sobel (1982): "Strategic information transmission," *Econometrica*, 50(6), pp. 1431–1451.

De los Santos, B., I. K. Kim, and D. Lubensky (2013): "Do MSRPs decrease prices?," *Working Paper*.

Diamond, P. A. (1971): "A model of price adjustment," *Journal of Economic Theory*, 3(2), 156 – 168.

Faber, R. P., and M. C. Janssen (2008): "On the effects of suggested prices in gasoline markets," Tinbergen Institute Discussion Papers 08-116/1, Tinbergen Institute.

Gill, D., and J. E. Thanassoulis (2010): "The optimal marketing mix of posted prices, discounts and bargaining," Economics Series Working Papers 479, University of Oxford, Department of Economics.

Haan, M. A., and J. L. Moraga-González (2011): "Advertising for attention in a consumer search model," *The Economic Journal*, 121(552), 552–579.

Horowitz, J. L. (1992): "The role of the list price in housing markets: theory and an econometric model," *Journal of Applied Econometrics*, 7(2), pp. 115–129.

Janssen, M., and S. Shelegia (2013): "Consumer search and double marginalization," Working paper.

Mathewson, G. F., and R. Winter (1984): "An economic theory of vertical restraints," *The RAND Journal of Economics*, 15(1), pp. 27–38.

Milgrom, P., and J. Roberts (1986): "Price and advertising signals of product quality," *Journal of Political Economy*, 94(4), pp. 796–821.

Puppe, C., and S. Rosenkranz (2011): "Why suggest non-binding retail prices?," *Economica*, 78(310), 317–329.

Rothschild, M. (1974): "Searching for the lowest price when the distribution of prices is unknown," *The Journal of Political Economy*, 82(4), pp. 689–711.

Tappata, M. (2009): "Rockets and feathers: understanding asymmetric pricing," *The RAND Journal of Economics*, 40(4), 673–687.

Telser, L. G. (1960): "Why should manufacturers want fair trade?," *Journal of Law and Economics*, 3, 86–105.

Thaler, R. (1985): "Mental accounting and consumer choice," *Marketing Science*, 4(3), pp. 199–214.

Wolinsky, A. (1986): "True monopolistic competition as a result of imperfect information," *The Quarterly Journal of Economics*, 101(3), pp. 493–512.

Yang, H., and L. Ye (2008): "Search with learning: understanding asymmetric price adjustments," *RAND Journal of Economics*, 39(2), 547–564.