Search and ripoff externalities

Armstrong, Mark

All Souls College, Oxford University

February 2015

Search and Ripoff Externalities

Mark Armstrong*

February 2015

Abstract

This paper surveys models of markets in which only some consumers are "savvy". I discuss when the presence of savvy consumers improves the deals available to all consumers in the market (the case of search externalities), and when the non-savvy fund generous deals for all consumers (ripoff externalities). I also discuss when the two groups of consumers have aligned or divergent views about market interventions. The analysis focusses on two kinds of models: (i) an indivisible product in a market with price dispersion, and (ii) products which involve add-on pricing.

Keywords: Add-on pricing, bounded rationality, consumer protection, consumer search, externalities, price dispersion.

1 Introduction

This paper examines situations in which "savvy" and "non-savvy" consumers interact in the marketplace. An old intuition in economics suggests that savvy consumers help to protect other consumers, and intervention to protect vulnerable consumers is needed only when there are insufficient numbers of savvy types present in the market. In broad terms, a "search externality" operates so that those consumers who are informed about the deals available in the market ensure that less informed consumers also obtain reasonable outcomes. More recent work, however, has examined situations where consumers benefit from the presence of non-savvy types. In such markets, a "ripoff externality" is present—some consumers end up buying services they do not value, say, which help fund generous deals elsewhere—and vulnerable consumers may need protection even when they are relatively few in number.

*Department of Economics and All Souls College, University of Oxford. Early versions of this paper were presented to the 2013 Annual Meeting of German Economic Association in Düsseldorf and to the 2014 conference on *Industrial Organization: theory, empirics and experiments* organized by the University of Salento. I am grateful for discussions on this topic to Kyle Bagwell, Renaud Foucart, Michael Grubb, Andrew Rhodes, Tom Ross, Rani Spiegler, Vic Tremblay, John Vickers, and Jidong Zhou.

This paper discusses two principal issues. First, what determines the direction of inter-consumer externalities in a market? That is, when do savvy consumers protect other consumers, when do non-savvy consumers improve the deals offered to the savvy, or when is there no interaction between the two groups at all? Second, which kinds of market interventions—such as policies to increase the number of sellers, to constrain high prices, or to increase the number of savvy types in the market—benefit both consumer groups and which benefit one group at the expense of the other?

For our purposes, there are two broad notions of savviness to consider. First, a consumer might be *well informed* about the prices, product qualities or her idiosyncratic value for the product. For instance, a savvy consumer knows whether a given wine will likely be to her taste just by looking at the label, or a savvy consumer looking for a new television may know the range of available prices (e.g., because she is online). Alternatively, a consumer might be *strategically* savvy, in that she has a good understanding of the game being played in the market. For instance, consumers might be unable to discern product quality (so they are not savvy in the first sense) but they understand how quality depends on price in equilibrium and buy accordingly. Or they might foresee a firm's incentive to set its future prices. A consumer who is savvy in this sense is aware of her future behaviour, while a strategically naive consumer might not predict accurately what she will want or need in the future.

A consumer might be non-savvy in both senses. For instance, she might not be able to discern quality and also might not foresee how quality depends on price. Indeed, strategic naivety might be the cause of information problems. In a market where in fact there is price dispersion but naive consumers think that all sellers offer the same price, for example, a naive consumer might choose not to incur the search cost required to become informed about prices in the market.

A framework for discussing these issues is the following.[1] Suppose there are two kinds of consumers, "savvy" and "non-savvy", and the proportion of savvy consumers in the population is σ. To focus on the impact of savviness on outcomes, I suppose that there are no systematic differences in *tastes* for the product in question across the two groups of consumers. For the most part, I take the extent of savviness, σ, to be exogenous and out of the control of consumers and sellers.

[1] For a similar approach, see Armstrong (2008, section III.C) and Armstrong and Vickers (2012).

Let $V_S(\sigma)$ and $V_N(\sigma)$ denote the expected net surplus enjoyed in equilibrium by an individual savvy and non-savvy consumer respectively, while $V(\sigma) \equiv \sigma V_S(\sigma)+(1-\sigma)V_N(\sigma)$ measures aggregate consumer surplus. We expect that $V_S(\sigma) \geq V_N(\sigma)$, so that savvy consumers obtain weakly better deals than their non-savvy counterparts. This is because tastes do not differ across the two groups of consumer, and a savvy type could mimic a non-savvy buying strategy and so obtain surplus V_N.[2] In situations where this makes economic sense, the difference $V_S(\sigma) - V_N(\sigma)$ represents a consumer's incentive to "become savvy" when σ other consumers are already savvy. A rational, but uninformed, buyer must obtain non-negative surplus $V_N \geq 0$, for otherwise she would choose to stay out of the market. However, a strategically naive consumer might experience negative surplus. In many cases V_S and V_N move the same way with σ—i.e., either both increase with σ, both decrease with σ, or neither depends on σ—although it is not inevitable this be so.[3]

Likewise, let $\Pi_S(\sigma)$ and $\Pi_N(\sigma)$ denote the profit generated in equilibrium by an individual savvy and non-savvy consumer respectively, while $\Pi(\sigma) \equiv \sigma\Pi_S(\sigma) + (1-\sigma)\Pi_N(\sigma)$ measures industry profit. Here, it is less clear how Π_S and Π_N compare, although in most of the situations discussed in this paper non-savvy consumers generate more profit than the savvy and $\Pi_S(\sigma) \leq \Pi_N(\sigma)$. In perfectly competitive situations we expect average profit to be zero, although profit extracted from one group might be used to subsidize the offer made to the other. Finally, let $W(\sigma) = V(\sigma) + \Pi(\sigma)$ denote total welfare when σ consumers are savvy.

In this paper I focus on three cases of interest:

Search externalities: When consumers are better off when the proportion of savvy types is larger—that is, when $V_N(\sigma)$ and $V_S(\sigma)$ increase with σ—I say that "search externalities" are present. This is because the leading example where savvy consumers protect non-savvy consumers is when the former are better informed about prices or qualities available in the market. When more consumers are aware of all the available deals this forces sellers to offer good deals, which in turn are available to non-savvy buyers (as well as to other savvy buyers).[4]

[2] However, as illustrated in section 2.3, there are situations in which replacing a population of savvy buyers with a population of non-savvy buyers will make buyers better off. There are also cases where the two kinds of consumer obtain the *same* surplus. (As in section 3.1, this is the case when all sellers offer the same single deal.)

[3] For instance, in section 3.3 it may be that V_N increases with σ while V_S decreases with σ.

[4] The headline of the UK's *Daily Telegraph* on 9 July 2014 was "Savvy shoppers force down prices".

Ripoff externalities: When individual consumers are better off when the proportion of savvy types is smaller—that is, when $V_N(\sigma)$ and $V_S(\sigma)$ decrease with σ—I say "ripoff externalities" are present. A leading example of this situation is when non-savvy consumers can be "ripped off" with extra charges, and the resulting revenue is passed back to all consumers in the form of subsidized headline price. It is possible that aggregate consumer surplus V rises with σ, even though both V_S and V_N fall with σ, if the gap $(V_S - V_N)$ is large (as is the case in section 3.3).

No interactions between consumers: On the boundary between these cases are situations in which there is no interaction between the two groups of consumers, and V_S and V_N do not depend on σ. These cases often involve biased beliefs on the part of non-savvy consumers. Here, competition delivers what each type of consumer thinks they want, and neither wishes to choose the deal offered to the other type. *Ex post*, though, biassed consumers might regret the deal they chose. (A lucky charm to help predict winning lottery numbers, say, has no impact on the savvy consumers who do not buy it, but may be attractive *ex ante* to gullible consumers.)

The plan for the rest of this paper is as follows. Models which generate price dispersion are examined in section 2, and I present three models to exemplify the three patterns of externalities listed above. Models with add-on pricing are presented in section 3, and again variants are chosen to illustrate the three patterns of externality. One lesson from the analysis is that small changes in model assumptions can swing the market from one kind of externality to the other, and the "small print" in the model matters. I end the paper with some concluding comments, including suggestions for markets outside these two families where search or ripoff externalities are likely to be present.

2 Price Dispersion

2.1 Search externalities

In a market for an indivisible good of known quality, it is plausible that when some consumers are aware of available prices and buy from the cheapest seller, those who shop less diligently are partially protected. To illustrate this, consider Varian (1980)'s classical model of price dispersion.[5] Here, n identical sellers supply a homogeneous product with

[5] See Salop and Stiglitz (1977) for closely related analysis. (This paper appears to have introduced the term "ripoff" into the academic economic literature.)

unit cost c. Consumers may differ in their reservation value for the item, v, where the fraction of consumers with $v \geq p$ is denoted $q(p)$. For ease of notation, write $\pi(p) \equiv (p-c)q(p)$ for profit with price p, which I assume is single-peaked in p, and p^M for the price which maximizes this profit. An exogenous fraction σ of consumers (independent of valuation v) are savvy, in the sense that they buy from the cheapest seller, while the other $1-\sigma$ consumers buy from a random seller so long as that seller's price is below their v.[6]

In cases where all consumers are savvy or all are non-savvy, there is a pure strategy equilibrium and no price dispersion. If $\sigma = 1$, so that all consumers shop around, there is Bertrand competition and price is driven down to cost c. If $\sigma = 0$, so that all consumers shop randomly, then no seller has an incentive to set price below the monopoly price p^M, and the outcome is as if a single seller supplied the market. Since there is no price dispersion, it follows that $V_N = V_S$ and $\Pi_N = \Pi_S$ in these extreme cases. (Here, V and Π refer to the expected value of a consumer's surplus and profit, with expectations taken over the idiosyncratic valuation v.)

However, in a mixed market with $0 < \sigma < 1$, the only (static) equilibrium involves a mixed strategy for prices, so there is price dispersion in the market and a savvy consumer obtains a (weakly) lower price than any non-savvy consumer. It follows that $V_S > V_N$ and $\Pi_S < \Pi_N$. In more detail, the symmetric equilibrium involves each of the n sellers choosing its price according to a cumulative distribution function (CDF) $F(p)$, which satisfies

$$\left[\sigma(1-F(p))^{n-1} + \tfrac{1}{n}(1-\sigma)\right]\pi(p) \equiv \tfrac{1}{n}(1-\sigma)\pi(p^M)\ . \tag{1}$$

Here, a seller which chooses price p will sell to all σ savvy consumers (who have $v \geq p$) provided all of its rivals choose a higher price, which occurs with probability $(1-F(p))^{n-1}$ in this equilibrium. On the other hand, the seller will always sell to its share of the $1-\sigma$ non-savvy consumers (who have $v \geq p$). As such, a seller's demand from the non-savvy consumers is less elastic than demand from the savvy. The left-hand side of (1) is therefore the seller's profit if it sets price p. Since the seller could decide only to serve its captive consumers, who are $\frac{1}{n}(1-\sigma)$ in number, with the monopoly price, the right-hand side represents a seller's available profit.[7] For a seller to be willing to play the mixed strategy

[6] This behaviour could be justified if each consumer's cost of search is very convex, in the sense that a consumer can visit one seller for free but finds it too costly ever to visit a second seller. A fraction σ are informed of each seller's price, while the remaining $1-\sigma$ consumers are informed of no price. An alternative interpretation of this inert behaviour is that $1-\sigma$ consumers are strategically naive, and mistakenly believe that competition ensures the "law of one price" operates and all sellers offer the same price.

[7] It is not an equilibrium for sellers to choose the monopoly price $p = p^M$ for sure, since a seller could

$F(\cdot)$, it must be indifferent between all prices in the support of $F(\cdot)$.

Given p, the value of $F(p)$ which solves (1) is an increasing function of σ. That is, when the fraction of savvy consumers is higher, each seller is more likely to set low prices. Intuitively, increasing σ expands the portion of demand which is elastic. Because each seller's price distribution is shifted downwards when σ rises, both the savvy (who pay the minimum price from n draws) and the non-savvy (who pay the price from a single draw) are better off when σ is higher. In the notation of section 1, then, V_S and V_N increase with σ, as does aggregate consumer surplus V. From (1), industry profit is $\Pi(\sigma) = (1-\sigma)\pi(p^M)$, which decreases with σ. Total welfare $W(\sigma)$ at least weakly increases with σ since lower prices stimulate demand.[8]

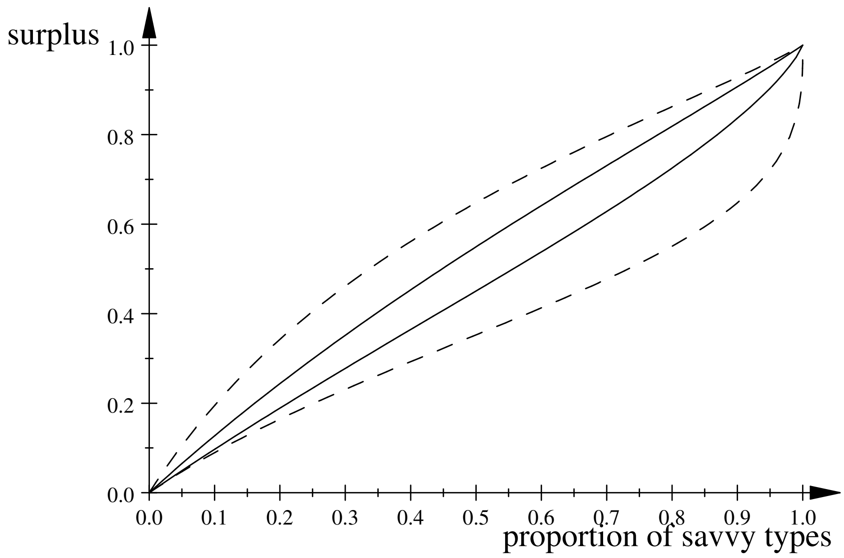

Figure 1: Expected surplus with price dispersion ($n = 2$ and $n = 4$)

Figure 1 depicts the net surpluses V_S and V_N enjoyed by the savvy (the upper solid curve) and the inert (lower solid curve) consumers when all consumers are willing to pay $v = 1$ for the item and $c = 0$, so that $\pi(p) = p$ if $p \leq 1$, and when $n = 2$. Note that the extent of price dispersion, as captured by the gap between the minimum and average price

slightly undercut this price and thereby serve all the savvy consumers.

[8] Armstrong and Zhou (2011, section 1) extend this model of price dispersion so that, instead of purchasing from a *random* seller, a sales intermediary steers the non-savvy consumers towards a supplier of his choice if given an incentive by that supplier to do so. (The savvy consumers are immune to the salesman's patter, observe the full list of retail prices, and buy from the cheapest supplier.) In equilibrium, there is a positive relationship between a supplier's choice of retail price and per-sale commission, and non-savvy consumers buy the most expensive product. One can show that the search externality is present in this market, and savvy consumers who are able to withstand sales pressure benefit all consumers in this market for advice.

in the market, is non-monotonic in σ. (As discussed, there is no price dispersion in the extreme cases where $\sigma = 0$ or $\sigma = 1$.) As such, increasing σ might increase or decrease the extent of price dispersion in a market, depending on the initial level of savviness.[9]

The two solid curves on Figure 1 are rather close together, indicating there is a limited benefit to a consumer in knowing both prices. When the number of sellers is larger, however, one can show that the expected price paid by savvy consumers falls while the expected price paid by the non-savvy consumers rises, so the two curves are stretched further apart. Intuitively, a firm's demand from the savvy consumers, $\sigma(1 - F)^{n-1}$, falls with n faster than its demand from the non-savvy, $(1 - \sigma)/n$, and so with larger n a firm puts more weight on extracting revenue from the latter group. (One can see that the prices paid by informed and uninformed consumers must move in opposite directions as n increases, since industry profit $\Pi(\sigma) = (1-\sigma)\pi(p^M)$ does not depend on n.) The dashed lines on the figure show the respective surplus functions in this example when $n = 4$. Thus, increasing the number of sellers has contrasting effects on the informed and the uninformed consumers, with the non-savvy suffering with "more competition" of this form.[10]

Making the fraction of savvy types endogenous: The discussion so far has taken the fraction of savvy types as exogenous. However, it may be that consumers can *choose* to be savvy, for instance by investing in acquiring market information. When information about market conditions and product attributes is costly to acquire, it may be rational to stay uninformed, especially when the search externality is present and most other consumers are already well informed.[11] To discuss the equilibrium extent of savviness, continue with Varian's model of price dispersion, and when the fraction of savvy types is σ write a savvy consumer's surplus as $V_S(\sigma)$ and the surplus of an uninformed consumer as $V_N(\sigma)$. (The following argument is easiest if I assume all consumers have the same value v for the product, so that all consumers will buy in equilibrium.) As illustrated on Figure 1, V_S and V_N increase with σ, while the incentive to become informed, $V_S(\cdot) - V_N(\cdot)$, is "hump-shaped" such that $V_S(0) - V_N(0) = V_S(1) - V_N(1) = 0$.

[9] Brown and Goolsbee (2002) find evidence consistent with this, when they observe that price dispersion rises when the use of price comparison websites increases from a low level, and then decreases as their use becomes more widespread.

[10] See Morgan, Orzen, and Sefton (2006) for further discussion of the impact of changing σ and n on payoffs to consumers. These authors also conduct an experiment, where human sellers face computer consumers, and which confirms the model's predictions quite closely.

[11] The issue of how many agents rationally decide to remain uninformed in equilibrium was highlighted early on by Grossman and Stiglitz (1980) and Burdett and Judd (1983).

Suppose that consumers can switch from being ignorant to informed by incurring an up-front information acquisition cost, κ. As discussed in Burdett and Judd (1983, section 3.2), a consumer will choose to become informed if and only if $V_S(\sigma) - V_N(\sigma) \geq \kappa$, and consumers will choose to become informed until the marginal consumer is indifferent. Thus, the fraction σ of consumers who become informed in an equilibrium with $0 < \sigma < 1$ satisfies

$$V_S(\sigma) - V_N(\sigma) = \kappa \ . \tag{2}$$

If κ is too large, there is no solution to (2) and the only equilibrium involves no consumers becoming informed and sellers charging the monopoly price.[12]

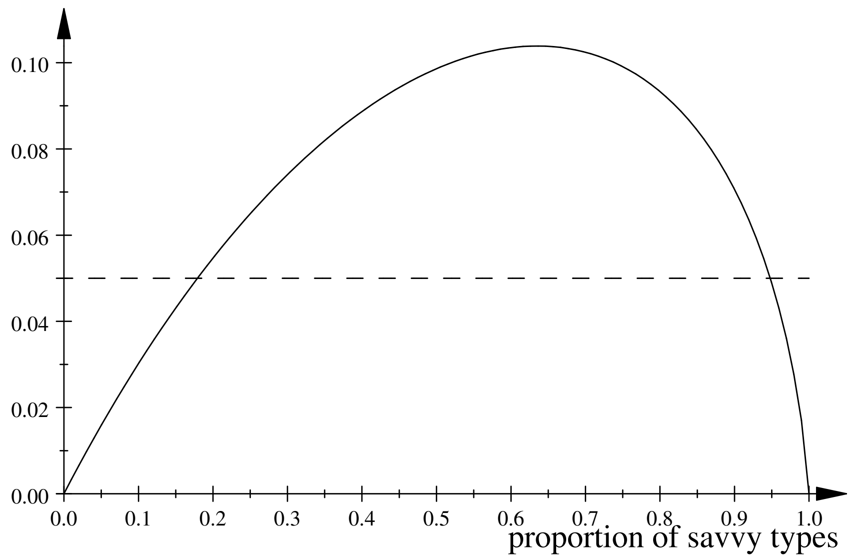

Figure 2: The fraction of consumers who choose to be informed

Figure 2 illustrates the incentive to become informed in the duopoly example from Figure 1, so that the hump-shaped curve depicts the difference between the two solid curves on Figure 1, $V_S(\sigma) - V_N(\sigma)$. The flat dashed line represents the cost of becoming informed (κ is equal to $\frac{1}{20}$ on the figure). The figure shows a situation with two interior equilibria satisfying (2). However, only the higher-σ equilibrium is stable, while at the low-σ equilibrium a perturbation in σ will induce σ to move away from this point. As emphasized by Grossman and Stiglitz (1980) in a related model, it is never an equilibrium for *all* consumers to become informed. In an interior equilibrium, the search externality

[12] Even if (2) has a solution, it is always one equilibrium for all consumers to remain uninformed. When no one is informed, all consumers obtain the same (bad) deal in the market, and there is no point in an individual consumer spending $\kappa > 0$ to acquire information about the available deals.

implies that too few consumers choose to be informed—too many prefer to free-ride on other consumers' search efforts—and consumer surplus would be boosted if σ were increased.[13]

One can imagine consumer policies which affect either left-hand or right-hand side of expression (2), i.e., either the cost or the benefit of being savvy. Assuming that it is the high-σ equilibrium on Figure 2 which is relevant, a policy which reduces information acquisition costs—so that κ is lowered—will increase σ, and this will in turn benefit all consumers. Likewise, a policy which shifts the benefit curve upwards will increase equilibrium σ. For example, we saw on Figure 1 that increasing the number of sellers pushed the surplus of the two groups of consumers further apart, and so shifted the benefit curve upwards. Since this will increase σ, it may be that increasing the number of sellers will benefit all consumers—not just the savvy—once the equilibrium impact on σ is taken into account.[14]

On the other hand, a policy which shifts the benefit curve downwards will reduce the fraction of consumers who choose to become informed.[15] Consider the situation with two sellers, consumer valuation $v = 1$ and costless production as depicted on Figures 1 and 2. Suppose that any consumer can become informed of both prices, rather than having to shop randomly, by incurring the cost $\kappa = \frac{1}{20}$. In this case, a fraction $\sigma \approx 0.95$ of consumers choose to be informed and all consumers have expected surplus (net of the search cost where relevant) of about 0.9, i.e., their outlay is around 0.1. Here, most consumers obtain what seems like a good deal, obtaining the item in return for a total outlay which is only 10% of their valuation. However, a few consumers will pay up to ten times this price, and pressure—from the media, politicians, or consumer groups—to protect consumers from these occasional high prices could arise. In response, suppose that a new policy constrains firms to set prices no higher than $\frac{1}{4}$, say, so that the maximum permitted price is one quarter of the highest price seen without regulation. For given σ, the expected prices

[13] Aggregate consumer surplus when σ consumers become informed is $\sigma(V_S(\sigma)-\kappa)+(1-\sigma)V_N(\sigma)$, which is strictly increasing in σ at any point satisfying (2).

[14] To take an extreme example, if all consumers have a relatively high information acquisition cost $\kappa = \frac{1}{5}$, then by examining Figures 1 or 2 we see that the only equilibrium with duopoly involves no consumers becoming informed, in which case all consumers are charged the monopoly price and obtain no surplus. However, with four suppliers, the maximum gap between V_S and V_N is greater than κ, and a stable equilibrium with $\sigma \approx 0.975$ emerges where all consumers obtain surplus (net of search cost where relevant) of about 0.78. A contrasting effect is discussed in Spiegler (2011, page 150): when a consumer is faced with a greater number of suppliers, she may suffer from "choice overload", with the result that *fewer* consumers are savvy.

[15] See Fershtman and Fishman (1994) and Armstrong, Vickers, and Zhou (2009) for analysis of this issue.

paid by the informed and uninformed consumers then quarter, and hence the incentive to become informed is also quartered. The result is that there is no solution to (2), and the only equilibrium involves all consumers remaining uninformed and sellers charging the maximum permitted price, $p = \frac{1}{4}$. Each consumer now has outlay of 0.25 rather than 0.1, while industry profit increases about five-fold increase with the new policy. Thus, the perverse effect of this policy on consumers can be substantial.[16]

So far, I have discussed how consumers can take the initiative to become savvy. Clearly, though, sellers also play a role in supplying information to consumers, and there is a vast literature about how seller advertise their product attributes and prices. Less familiar is the possibility that sellers attempt to "confuse" consumers, with the result that the fraction of savvy types falls. For example, sellers might present their prices in an opaque way or in a format different to their rivals, and this makes it hard for consumers to compare deals.[17]

To illustrate this possibility, consider the following extension to Varian's model.[18] There are two sellers, and a seller can present its price in one of two formats. (A seller chooses its price and its format simultaneously.) If sellers choose the same format, consumers find it easy to compare prices and all of them choose to buy from the seller with the lower price. However, if sellers choose distinct formats a fraction $1 - \sigma$ of consumers are confused and buy randomly (while the remaining σ are savvy enough to make an accurate comparison

[16] Knittel and Stango (2003) examine the credit card market in the United States in the period 1979–89, during which usury laws in some states put a ceiling on permitted interest rates. In their Table 3 they show how, for much of this period, average interest rates were *higher* in those states with a ceiling, and interpret this as evidence that price caps can encourage tacit collusion via a policy-induced focal point. The (static) search model presented in the text provides an alternative explanation for why a price cap might lead to high prices, although the particular example discussed would look as if the sellers were coordinating on the price cap.

[17] Several papers provide evidence of consumer confusion due to formatting problems. For instance, Clerides and Courty (2013) observe empirically that the same brand of detergent is sold in two sizes, the large size containing twice as much as the smaller. Sometimes the large size is more than twice as expensive as the smaller, and yet significant numbers of consumers still buy it. An example of a tariff which may be difficult for some consumers to compare with rival tariffs is described in the *Royal Mail*'s "handy guide" to its prices—see www.royalmail.com/sites/default/files/RM_OurPrices_Dec2014.pdf [visited 26 January 2015]. This determines the price for delivering a parcel as a function of four physical characteristics (length, width, depth and weight). For instance, a specific price applies for international delivery for a parcel with "Height + Width + Depth no greater than 90cm with no single side longer than 60cm, Weight up to 2kg".

[18] This discussion is based on Piccione and Spiegler (2012) and Chioveanu and Zhou (2013). Wilson (2010) takes a different approach. In his model, two sellers compete to supply a homogenous product, and each seller can choose the cost that consumers incur to discover its price. A fraction of consumers can understand prices without cost, regardless of seller strategies, while the search cost of the remaining consumers is affected by obfuscation strategies. (All consumers can observe each seller's obfuscation strategy from the start.) Wilson shows that in equilibrium one seller chooses a high search cost in order to relax subsequent price competition.

even across formats).

In this context, sellers choose both price and format according to a mixed strategy. Since the format itself does not matter, only whether formats are the same or not, a seller chooses the same CDF for its price, say $F(p)$, regardless of its chosen format, and is equally likely to choose either format. If a seller chooses a particular format and price p, then as in expression (1) its expected profit is

$$\left(\underbrace{\tfrac{1}{2}[1-F(p)]}_{\text{same format}}+\underbrace{\tfrac{1}{2}[\sigma(1-F(p))+\tfrac{1}{2}(1-\sigma)]}_{\text{different format}}\right)\pi(p)\equiv\tfrac{1}{4}(1-\sigma)\pi(p^M)\ .$$

Here, if the two sellers display their prices in the same format there is fierce competition, and the cheaper seller wins the whole market, while if the formats differ a fraction $(1-\sigma)$ of consumers shop randomly. The right-hand side of the above represents the profit obtained when a seller uses a different format and fully exploits its captive consumers, which is each seller's equilibrium expected profit.

It is not an equilibrium in this model for sellers to choose their format deterministically. Clearly, if both sellers chose the same format for sure, price would be driven down to cost and profit to zero. In that case, a seller could switch format to make money from the newly confused consumers. If sellers were known to choose distinct formats, prices would be chosen according to a mixed strategy as in (1). However, in that case a seller could switch to offer the same format as its rival and offer the lowest price in the price support, which ensures it serves the entire market and boosts its profit.[19]

This model predicts that sellers engage in "tariff differentiation" to obtain positive profit, just as sellers in more traditional oligopoly models engage in product differentiation. However, unlike forms of product differentiation, this tariff differentiation confers no welfare gains. A consumer policy which forced firms to present prices in a common format would, in this stylized model, lead to Bertrand price competition, and all consumers would benefit.[20]

[19] Clearly, if sellers choose their format first, and subsequently choose price, in this model the most profitable subgame-perfect equilibrium involves the two sellers choose different formats for sure, since whenever they have the same format their profits will fall to zero.

[20] As discussed by Piccione and Spiegler (2012) and Chioveanu and Zhou (2013), additional issues arise when the two formats are "simple" and "opaque", and when both sellers choose an opaque format even more consumers are confused relative to when sellers choose distinct formats. In such a setting, when a seller sets a low price it chooses a simple format to make it easy for customers to see its low price, but with a high price it offers an opaque format. Alternatively, it may be that when consumers find it hard to compare the sellers' offers they exit the market instead of buying randomly. If so, this could discipline unregulated sellers to present their offers in a comparable format. (See Crosetto and Gaudeul (2014) for an experimental test of this possibility.)

2.2 No externalities

Varian's model of price dispersion can be extended so that sellers are able to charge distinct prices to savvy and non-savvy consumers. For example, the former group might be those who use a price-comparison website and buy online, while the uninformed go to a random bricks-and-mortar store, and a seller might set different prices for the two purchase channels. When this form of price discrimination is used, the link between the two groups is broken, and the outcome is that the informed consumers are offered a low price equal to marginal cost c, while the uninformed pay the monopoly price p^M. In this case, there is no search externality and the fraction of informed consumers has no impact on the surplus enjoyed by a consumer of either type.[21]

Somewhat related is the possibility that sellers might each offer several tariffs: savvy consumers pick the lowest price from all tariffs from all sellers, while non-savvy consumers buy more randomly from the set of available tariffs. If a seller offers one cheap tariff alongside many expensive options, it is in effect able to compete separately for the two groups of consumers, with only the savvy consumers obtaining good deals.[22]

A less obvious way to eliminate the search externality is to place the workhorse model of section 2.1 in a dynamic context, and to consider the impact of consumer savviness on the sustainability of tacit collusion.[23] Suppose the industry attempts to collude at the monopoly price p^M with the use of a trigger strategy. If a seller deviates by undercutting p^M, this is detected by all rivals, and from the next period onwards the industry plays the one-shot Nash equilibrium with mixed strategies described above, yielding per-seller profit in each period given by the right-hand side of (1). Crucially, if a seller does undercut the collusive price, only the σ savvy consumers can react. As a result, when δ is the discount

[21] Baye and Morgan (2002) consider a model in which sellers must pay to list on a price comparison website, and can charge different prices on this website and when they sell direct to consumers. They find that sellers choose whether to list, and their price on the comparison website, according to a mixed strategy and choose their price on the comparison website according to a mixed strategy. Sellers obtain positive profits there because of the possibility they are the sole listing seller. A seller's price on the comparison website is lower than its price on its own platform.

[22] This is similar to the previous discussion about confusing tariffs, except that here different customers of the same seller can pay different prices for the product. Miravete (2013) documents when a seller offers a tariff which is dominated by other tariffs it offers—which he terms "foggy pricing"—in mobile telephony.

[23] See Schultz (2005) for this analysis, as well as its extension to a market with horizontally differentiated products. Petrikaite (2014) analyzes an alternative model in which consumers become informed about prices and valuations by searching sequentially through their options. She finds that an increase in the cost of search—i.e., a reduction in market transparency—usually makes collusion easier to achieve.

factor, collusion at the monopoly price can be sustained if

$$\underbrace{\frac{1}{1-\delta}\frac{\pi(p^M)}{n}}_{\text{collusive profit}} \geq \underbrace{(\sigma + \frac{1-\sigma}{n})\pi(p^M)}_{\text{deviation profit}} + \underbrace{\frac{\delta}{1-\delta}(1-\sigma)\frac{\pi(p^M)}{n}}_{\text{punishment profit}}$$

which reduces to the familiar condition

$$\delta \geq \frac{n-1}{n} \ . \tag{3}$$

In this market, an increase in σ has two contrasting effects. When σ is large there is fierce competition without collusion, and so the punishment profit is low. On the other hand, when σ is large, the number of consumers who are able to respond to a price cut is large, and so the short-run gains from deviating are large. These countervailing effects precisely cancel out, and the ability to collude is unaffected by the number of savvy consumers. In particular, when the discount factor is large enough that (3) holds, all consumers pay the monopoly price p^M, and savvy types cannot protect consumers from exploitation.

2.3 Ripoff externalities

In this section, I consider one natural way to reverse the search externality in a market with price dispersion. In the workhorse model in section 2.1, price dispersion arose in equilibrium because a seller had some consumers who were captive and some who were able to shop around. This situation can also arise when all consumers see all prices but some consumers only find a single seller's product *suitable*. In this kind of market with horizontal product differentiation, sellers face a trade-off between exploiting their captive consumers and attracting those who like more than one seller's product. In this context, suppose that non-savvy consumers are unable to discern which products are suitable for them until after purchase. These consumers view products as perfect substitutes *ex ante*, and so act to intensify price competition, to the benefit of all consumers.[24]

In more detail, suppose that two symmetric sellers costlessly serve a market. With probability α, where $0 < \alpha < 1$, a consumer finds a given seller's product to be suitable,

[24] Anderson and Renault (2000) study a related model with costly sequential search, where there is no price dispersion in equilibrium. Consumers have idiosyncratic tastes for the sellers' products. Savvy consumers know their tastes in advance, and travel to the seller with the preferred product first, while other consumers must travel to a seller to discover their match utility. (All consumers must travel to a seller to discover its price.) Here, as in the model presented in the text, informed consumers have less elastic demand and their presence boosts industry profits and harms other consumers.

in which case she values its product at v. With probability $1-\alpha$ the product is unsuitable, in which case it is worthless to her. All consumers can see both sellers' prices. A fraction σ of consumers are savvy in the sense that they can discern their match utility in advance, and buy from the cheapest seller with a good match (if any). A fraction $1-\sigma$ of consumers cannot judge the match quality until after they have purchased the product. These consumers are rational and risk-neutral, and buy from the cheapest seller (if any) with a price below the expected match quality, which is αv. Of course, these non-savvy consumers end up with a worthless product with probability $1-\alpha$, and in this sense they are "ripped off". However, unlike the final model in this paper, this feature is not due to sellers engaging in any tactic which aims to exploit the non-savvy consumers.[25]

It is clear that any equilibrium involves sellers choosing price according to a mixed strategy, and in the appendix I derive the symmetric mixed strategy equilibrium in this market.[26] There, we see that industry profit is

$$\Pi(\sigma) = 2\sigma v\alpha(1-\alpha) \ , \tag{4}$$

which *increases* with the fraction of savvy types, σ. I also show that, as σ increases, a seller's price weakly increases in the sense of first-order stochastic dominance. This implies that a consumer of either type is better off when the fraction of non-savvy types rises, so that the rip-off externality is present. In the case where $v = 2$ and $\alpha = \frac{1}{2}$, $V_N(\sigma)$ is plotted as the lower solid curve on Figure 3, while $V_S(\sigma)$ is plotted as the upper solid curve. Aggregate consumer surplus, $\sigma V_S(\sigma) + (1-\sigma)V_N(\sigma)$, is plotted as the dashed curve, which is also decreasing (except for σ very close to 1).

Total welfare depends in a complicated way on σ in this model, with two conflicting effects: (i) increasing σ implies that more consumers are able to choose a suitable product, and (ii) decreasing σ means that more consumers treat products as homogeneous, which drives down prices and makes the non-savvy more likely to buy at all. (In the example where $v = 2$ and $\alpha = \frac{1}{2}$, a non-savvy consumer is never offered a low enough price to induce her to buy when $\sigma = 1$.) Because of these two effects, it is possible for welfare to

[25] One could extend this model so that sellers could choose their α parameter (say, where a larger α requires a higher fixed cost), in which case choosing a low α might be interpreted as an attempt to "rip off" non-savvy consumers who cannot discern their match utility.

[26] The analysis is a little more involved than that for section 2.1, since the two groups of consumers have different reservation prices: a savvy consumer is willing to pay up to v for a suitable product, while a non-savvy consumer is willing to pay only αv for any product. An implication of this is that the support for prices might have a "gap", and intermediate prices are never offered by sellers.

be non-monotonic with σ, although it is always maximized at $\sigma = 1$, when all consumers buy a product they like when one such product is available.

Note that the gap between V_S and V_N on Figure 3 is approximately 0.5 for all σ. Consider an extension to this model where any consumer can choose to become informed about their match utilities by incurring an *ex ante* cost of $\kappa = 0.1$, say. A rational consumer will therefore choose to be informed, regardless of how many others do so. When all consumers choose to become informed, however, the figure shows they are worse off compared to the situation where all remain ignorant. Here, the ripoff externality implies that too *many* consumers choose to become informed.

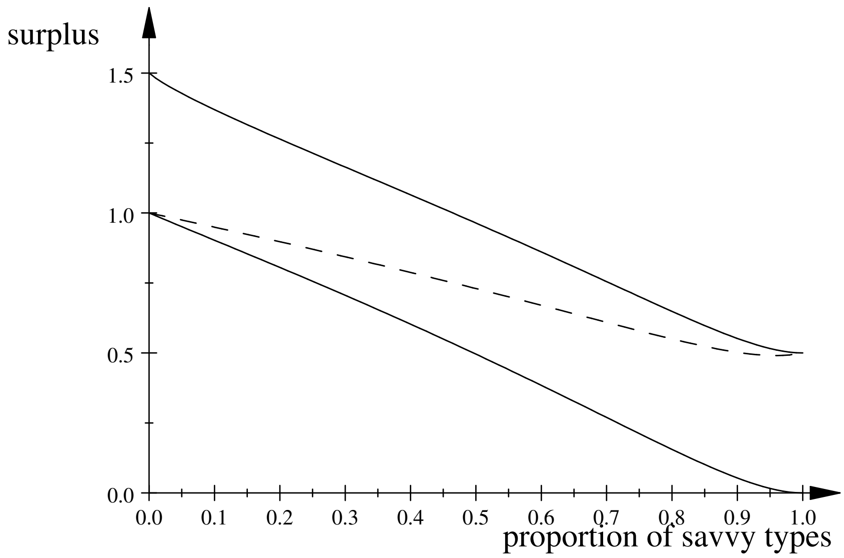

Figure 3: Expected surplus for the two groups of consumers

A final observation is that this model can be reinterpreted so that the roles of savvy and non-savvy consumers are reversed. Suppose that marketing efforts by these sellers can induce non-savvy consumers to view the two products as differentiated (some of whom then value the product more than it is really worth, and some of whom mistakenly view it as worthless), while savvy consumers are immune to this marketing and correctly view the products as perfect substitutes. With this reinterpretation, the search externality is present and savvy consumers protect the non-savvy.

3 Add-On Pricing and Aftermarkets

In this section I examine more complicated products than those discussed in section 2. Specifically, I discuss markets in which sellers supply an "add-on" product or service once a consumer has purchased an initial "core" product. Familiar examples of this phenomenon include: the minibar inside a hotel room; toner cartridges once one has purchased a printer; after-sales care for your new car; an extended warranty for your new television; renewing an initial magazine subscription; a casual overdraft from your bank, or the ability to have your luggage stowed in the aircraft's hold in the event it is deemed too large for the cabin. In such markets, non-savvy consumers might downplay the importance of add-on terms, or cannot easily observe or interpret such terms, when they decide on their supplier.

I focus on situations where a seller chooses its core and add-on prices at the same time, and the issue is not one of lack of commitment to a future add-on price. Rather, some consumers either do not observe the firm's choice of add-on price, or can observe it but do not think it will apply to them.[27] Three variants are discussed in turn: one where duopolists compete in a Hotelling market and must offer the same add-on terms to all their customers; a second where potentially many sellers compete and naive consumers do not foresee their demand for the add-on service, and a final variant where non-savvy consumers can be tricked into paying for add-ons they don't want. These model variants are chosen to illustrate the three kinds of externality listed in section 1.

Although the following discussion is couched in terms of add-on pricing, it largely applies also to markets where sellers choose the *quality* of an indivisible product, and where only some consumers are able to discern quality before purchase. (Here, a low add-on price corresponds to a high-quality product.) For instance, some consumers may not know how to interpret "exclusions" in an insurance contract, or may mistakenly believe that such exclusions will not be relevant to their circumstances.

3.1 Search externalities

Consider the following Hotelling duopoly model of add-on pricing. Two symmetric sellers, denoted 1 and 2, compete to supply a product. These sellers are located at each end of the unit interval $[0, 1]$, and consumers are uniformly located on this interval. A seller offers a

[27] Shapiro (1995) discusses four potential sources of market failure in aftermarkets, including lack of ability to commit to future prices and the presence of consumers who are poorly informed or myopic about future prices. My focus in this section is on the latter possibility.

core product, and if a consumer buys its core product she is then able to buy an add-on product from the same seller. All consumers want a single unit of the core product, and if the price for the add-on product is p they will go on to purchase $q(p)$ units of that product.[28] The net surplus to a consumer from the option of being able to buy the add-on at price p is denoted $s(p)$, the usual area under the demand curve:

$$s(p) \equiv \int_p^\infty q(\tilde{p})d\tilde{p} \ .$$

For technical reasons, suppose that $q(p)$ is logconcave in p. Write $\pi(p) \equiv (p-c)q(p)$ for the add-on profit with price p, and p^M for the price which maximizes this profit. (Since $q(\cdot)$ is logconcave, $\pi(\cdot)$ is single-peaked in p.) Each seller incurs unit cost for supplying the core product and the add-on product equal respectively to C and c.

In this section, I assume that a seller must offer the same add-on price to all its customers. Suppose that seller $i = 1, 2$ chooses price for the core product and its add-on product equal respectively to P_i and p_i. A consumer located at $\ell \in [0, 1]$ obtains net surplus

$$X - P_1 + s(p_1) - t\ell$$

if she buys from seller 1, where $t > 0$ is the "transport cost" involved in travelling a unit distance to the seller. (Here, X is each consumer's value for the core product, and so her total surplus from the core product at price P_1 plus the option of being able to buy the add-on at price p_1 is $X - P_1 + s(p_1)$.) Likewise, if this consumer buys from seller 2 her net surplus is

$$X - P_2 + s(p_2) - t(1-\ell) \ .$$

To simplify the analysis, I assume that X is large enough that the market is covered, and all consumers buy from one seller or the other (even if they foresee monopoly prices in the add-on market).

All consumers observe both sellers' core price. Suppose a fraction σ of consumers observe each seller's add-on price, while the remaining $1 - \sigma$ consumers either cannot observe, or do not consider, a seller's add-on price until after they have purchased their

[28] This elastic demand for the add-on service could be generated if each consumer has a unit demand for the add-on with incremental valuation v, and the probability that v is above p is $q(p)$. With this interpretation, the realization of v is not known to the consumer (even a savvy consumer) until after she buys the core product.

core product, at which point they are locked into the same seller for the add-on. I assume that whether a consumer is savvy is independent of her location ℓ.

A particularly simple model to analyze has the $1-\sigma$ non-savvy consumers being strategically naive, in that they can observe each seller's add-on price but at the time they purchase the core product they mistakenly believe they will have no demand for the add-on service. These consumers therefore care only about core prices when choosing their initial supplier, while savvy consumers care about the "lifetime" cost of the product they purchase.

In more detail, a seller's total profit if it chooses the price pair (P, p), while its rival chooses the equilibrium price pair (P^*, p^*) is

$$\left[\sigma\left(\frac{1}{2}+\frac{[s(p)-P]-[s(p^*)-P^*]}{2t}\right)+(1-\sigma)\left(\frac{1}{2}+\frac{P^*-P}{2t}\right)\right]\times[P-C+\pi(p)]\ . \quad (5)$$

To understand this expression, note that $[P-C+\pi(p)]$ is the seller's total profit from each of its customers, which consists of profit from the core product, $P-C$, plus the profit from the add-on, $\pi(p)$. The non-savvy do not react to changes in the add-on price, but do react to deviations in the core price P. Given that ℓ is uniformly distributed on $[0,1]$ and all consumers are served, the fraction of non-savvy consumers who prefer this seller's offer is $\left(\frac{1}{2}+\frac{P^*-P}{2t}\right)$. Savvy consumers react to deviations in both prices, which explains the extra term $s(p)-s(p^*)$ in the corresponding expression for their demand.

For prices (P^*, p^*) to constitute an equilibrium, choosing $(P, p) = (P^*, p^*)$ should maximize (5) over any pair of prices (P, p). The first-order conditions for this problem are

$$P^*-C+\pi(p^*)=t \quad (6)$$

and $\pi'(p^*)+\sigma s'(p^*)=0$ or

$$(1-\sigma)q(p^*)+(p^*-c)q'(p^*)=0\ . \quad (7)$$

Since $q(\cdot)$ is logconcave, there is a unique solution to (7) and hence to (6). It is harder to check whether the second-order condition for maximizing (5) is satisfied, but in standard examples this appears not to be an issue.[29] Note that the add-on price in (7) does not depend on differentiation parameter t, indicating that market power in the add-on market in this model depends on the extent of consumer myopia rather than horizontal product differentiation.

[29] For instance, if $q(p)=1-p$ and $C=c=0$, the second-order condition is satisfied for all σ and t.

Expression (6) implies that industry profit in equilibrium, $P^* - C + \pi(p^*)$, does not depend on the fraction of savvy types, σ. This is a common feature in Hotelling models with full coverage, although, as we will see shortly, the corresponding model with rational consumers has profit which does depend on σ. However, σ does affect how that profit is generated, as can be seen in expression (7). In particular, when $\sigma = 1$, the add-on price is at its efficient level $p^* = c$, while when $\sigma = 0$ the add-on price is the monopoly price p^M which maximizes $\pi(p)$. In the latter case, although a seller makes high profit $\pi(p^M)$ in the add-on market, this anticipated profit makes it keen to attract customers with a low core product price, and the net impact on profit is the same as when all consumers are savvy. More generally, formula (7) implies that the add-on price decreases monotonically from the monopoly to the efficient level as σ increases from 0 to 1. For example, when $q(p) = 1 - p$ and $c = 0$, the add-on price is $p^* = \frac{1-\sigma}{2-\sigma}$.

Total welfare rises when the add-on price moves closer to its efficient level, i.e., when σ is higher, and since profit is unaffected we deduce that aggregate consumer surplus, $V(\sigma) = V_N(\sigma) = V_S(\sigma)$, also rises with σ. Even if a consumer finds it prohibitively costly to discover the add-on price, when enough others do check these terms she is still able to obtain a reasonable deal. Thus, savvy types help to protect the non-savvy and this market exhibits search externalities.[30]

A natural question to ask is how this analysis is affected if the $1-\sigma$ non-savvy consumers cannot observe or interpret the add-on price but are rational instead of naive, and correctly forecast their own future demand for the add-on service as well as a seller's incentive to choose its add-on price. When rational consumers see one dimension of a seller's choice but not another, the issue arises of how the consumer forms her expectation of the unobserved variable given what she does observe. Here, I suppose that if a seller offers a particular core price (not necessarily the equilibrium price), the uninformed consumer calculates the seller's most profitable choice of add-on price given its core price, and chooses whether to buy from the seller accordingly.[31]

[30] As Shapiro (1995, page 493) puts it: "Poorly informed buyers may be protected by informed buyers, whose presence forces sellers to compete on a [total cost of ownership] basis and penalizes sellers with high aftermarket charges, especially since it may be difficult for sellers to identify the poorly informed buyers so as to price discriminate against them."

[31] These are known as "wary beliefs", following McAfee and Schwartz (1994). By contrast, "passive beliefs" involve a consumer's anticipated add-on price not depending on the seller's choice of core price. Passive beliefs are perhaps less plausible in this context, since if a seller deviates in its core price and rational-but-uninformed consumers do not change their beliefs about the add-one price, the seller in fact has an incentive to deviate in both prices. Nevertheless, it is considerably easier to solve models with

When $\sigma = 0$ or $\sigma = 1$ the outcome is exactly as the previous model with naive consumers, and in particular the industry profit is t in either of the extreme cases.[32] In a mixed market with $0 < \sigma < 1$, though, the two models differ, and consumer beliefs about the unobserved add-on price play an important role. In fact, when some savvy consumers are present, an rational-but-uninformed consumer anticipates that a seller with an unexpectedly low core price has set a *higher* add-on charge: when its core price is lower, the seller has less incentive to attract more custom from the savvy consumers with a low add-on price, and hence finds it profitable to set a higher add-on charge. This implies that a seller's demand is less elastic with respect to its core price than would be the case in a situation where all consumers were informed—or all were uninformed—and competition is less intense.[33] It follows that industry profit is higher in a mixed market than in a market with $\sigma = 0$ or $\sigma = 1$. For the same reason, consumers are worse off with intermediate σ than at the two extremes, so that $V = V_N = V_S$ is "U-shaped". Thus, consumers might be made better or worse off as σ is boosted—i.e., there might be a search or a ripoff externality—depending on the initial proportion of savvy types.

	$\sigma = 0$	$\sigma = 0.25$	$\sigma = 0.5$	$\sigma = 0.75$	$\sigma = 1$
p^*	0.50	0.42	0.29	0.12	0.00
P^*	0.75	0.83	0.96	1.05	1.00
$P^* - s(p^*)$	0.625	0.66	0.71	0.66	0.50
Π	1.00	1.08	1.17	1.16	1.00

Table 1: Outcomes in example with $q(p) = 1 - p$, $C = c = 0$ and $t = 1$

Solving this model is relatively complex, and I leave the details of the analysis to the appendix. Table 1 presents the outcome in an example where $q(p) = 1 - p$, $C = c = 0$ and $t = 1$. As can be seen, the equilibrium add-on price falls monotonically with σ, starting

passive than wary beliefs. (In this context, if the rational consumers had passive beliefs, the outcome is exactly as just described in the Hotelling model with naive consumers.)

[32] Ellison (2005) presents an alternative Hotelling model of add-on pricing, where consumers vary in both their the transport cost parameter t and their value for the add-on product. (Consumers with higher value for the add-on have stronger brand preferences.) He analyzes two games: one where the two firms reveal both of their prices *ex ante* and another where neither firm reveals its add-on price until consumers buy the core product. Using the current notation, these two cases correspond to situations with $\sigma = 1$ and $\sigma = 0$ respectively. In his model, industry profits are higher when no consumer is informed of add-on prices, in contrast to the model presented in the text.

[33] Scitovsky (1950, page 50) makes this point long ago: "the ignorant buyer's habit of judging quality by price weakens [...] price competition. [...] In such markets a price change will lead few buyers to transfer their customer from one producer to another. Hence, the price elasticity of demand will be low in such markets."

at the monopoly price p^M when $\sigma = 0$ and ending at the efficient price when $\sigma = 1$. The "total price", $P^* - s(p^*)$, which is inversely related to consumer surplus, is hump-shaped in σ, as is industry profit. Profit is the same at the two extremes $\sigma = 0, 1$, but is higher in a mixed population.

The ripoff externality is present when σ is small since when no consumers see the add-on price, they anticipate monopoly terms for sure and react sensitively to changes in a seller's core price. When some savvy consumers are introduced to the market, though, this induces rational-but-uninformed consumers to infer a high add-on price when they observe a low core price, and makes their demand less elastic and allows sellers to raise their core prices. Nevertheless, when only few non-savvy types are present, they are well protected in their purchases by the presence of savvy, well-informed buyers.

3.2 No externalities

An important reason why there was a search externality in the previous section is that consumers were assumed to be locked into their initial seller for the supply of the add-on service. If an undifferentiated add-on service was available from all sellers, competition for the core product is unaffected by profit in the aftermarket (which is zero), and there are no externalities across consumers.[34]

Another important reason for the search externality was the assumption that a seller had to offer the same add-on terms to all its customers, which seems reasonable in many contexts. (It is hard to imagine a hotel supplying rooms with different minibar prices, for instance.) In other situations, though, it may be feasible and profitable for sellers to offer distinct contracts aimed at the two groups of consumers.

To consider this second point in more detail, consider a situation in which non-savvy types do not anticipate their future demand for the add-on service. These consumers might observe the add-on price, but do not regard it as relevant to them. In the previous Hotelling model, suppose now that a seller can offer two contracts, and seller $i = 1, 2$ offers a price pair (P_i^S, p_i^S) aimed at the savvy type and the price pair (P_i^N, p_i^N) aimed at the naive. Suppose, hypothetically, that sellers can actually observe directly whether or not a consumer is savvy or naive, and can condition their tariff on the consumer type. (We will

[34] For example, some people may underestimate the costs of owning a dog (in terms of dog food, say), while others accurately forecast such costs. If the market for dog food is competitive, the price for a dog is not subsidized by aftermarket profits, and does not depend on the fraction of savvy types in the market.

see shortly that consumers will voluntarily choose the contract aimed at them, and so the implausible assumption that savviness is observable is not needed.)

It is easy to see that the equilibrium contract aimed at the savvy types is as in (6)–(7) with $\sigma = 1$, so that

$$P^S = C + t \ ; \ p^S = c \ , \tag{8}$$

while the contract aimed at the naive corresponds to the case with $\sigma = 0$, so that

$$P^N = C + t - \pi(p^M) \ ; \ p^N = p^M \ . \tag{9}$$

Thus, savvy types are offered a cost-reflective tariff, while naive types are offered a "bargain" core price but end up paying monopoly prices for the add-on once they realize they need or want this service. *Ex ante*, a naive consumer does not prefer the savvy contract, since the latter contract involves a higher core product price, and these consumers only care about that price at the time they choose seller. In addition, a savvy consumer does not prefer the "bargain-then-ripoff" contract aimed at the naive.[35] Thus, each consumer chooses the appropriate contract, and sellers do not need to observe savviness directly to discriminate between consumers. This model predicts, for instance, that a seller of cars might offer a menu of contracts to its customers: a bargain price for the car only, without aftercare services bundled in (which would then be expensive), aimed at naive consumers who do not foresee the importance of aftercare, and a bundled contract for the car plus specified aftercare, aimed at savvy consumers who care about the lifetime costs of the product.

There is therefore price dispersion in the market, both for the core product and for the add-on. The pair of contracts (8)–(9) does not depend on the proportion of savvy types, σ, and so there are no externalities between the two groups. In particular, the search externality found in section 3.1 was contingent on an assumption that the duopolists offered the same contract to all their customers.

Additional effects emerge in a modified model in which market participation is elastic rather than fixed as in the Hotelling framework. To that end, consider a perfectly competitive market with several identical sellers supplying an undifferentiated product.[36]

[35] Ignoring transport costs, a savvy consumer's net surplus with the savvy contract is $s(c) - (C + t)$, while her net surplus with the naive contract is $s(p^M) - (C + t - \pi(p^M))$, and the latter is smaller since add-on welfare $s(\cdot) + \pi(\cdot)$ is maximized at $p = c$.

[36] We can think of sellers each offering a menu of contracts, or separate sellers each offering a single contract. In the latter case we require at least four sellers, so that there are at least two sellers which offer savvy contracts and two which offer naive contracts.

Consumers differ in their valuation for the core product, X, and the fraction of consumers with $X \geq P$ is denoted $Q(P)$. Regardless of their valuation, X, all consumers will have the same add-on demand $q(p)$. Savvy consumers buy the product if their lifetime utility exceeds the core product's price, i.e., if $X + s(p) \geq P$. There are two natural ways to model the naive consumers' purchase decision.[37]

Case 1: Hidden costs. Here, naive consumers do not realize they need the add-on service, and when they buy the core product they in effect think the complementary product comes for free. For example, (very) naive consumers may not realize that toner cartridges are needed to use a printer, and so buy a printer assuming they can print as much as they wish without further outlay, or an inexperienced driver does not anticipate that adequate servicing in required to keep her car on the road. These consumers behave as if the add-on price will be zero, and so buy the core product if $X + s(0) \geq P$; they over-estimate the combined benefit of the core and add-on product and purchase too often.

Case 2: Hidden benefits. Naive consumers do not realize they will gain any benefit from the add-on service until after they buy the core product. For example, a naive consumer when choosing a hotel room overlooks the benefits of having the minibar in the room. The consumers behave as if the add-on price will be infinite, and so buy the core product if $X \geq P$; they under-estimate the combined benefit of the system and purchase too rarely.

In either case, once they have purchased the core product, a naive consumer goes on to generate profit $\pi(p)$ for the seller. Similarly to the contracts (8)–(9), the equilibrium contracts take the form whereby savvy consumers have the efficient contract $(P, p) = (C, c)$ while naive consumers have "bargain-then-ripoff" contract with a monopoly add-on price p^M and a subsidized core product price which just enables a seller to break even, so that $P - C + \pi(p^M) = 0$. As before, consumers have no incentive to choose the contract aimed at the other type, and this pair of contracts does not depend on the fraction of savvy types present in the market.

This model predicts that naive consumers end up paying high add-on prices. As such, regulators might consider controlling a seller's freedom to exploit naive consumers in this fashion.[38] (Such regulation has no impact on the surplus enjoyed by savvy consumers.) With the situation in case 1 above, a policy which forces sells to set $p = c$, say, will benefit

[37]This distinction did not matter with the previous Hotelling model, since naive consumers purchased from one seller or the other regardless of their anticipated surplus from the aftermarket.

[38]See Grubb (2015b) for an account of policy interventions when consumers have biased beliefs.

the naive consumers: it improves their surplus in the aftermarket, and the resulting price rise for the core product mitigates their problem of purchasing the combined system too frequently. However, with case 2 the two market failures in the *laissez-faire* market work against each other—naive consumers pay too much for the add-on, and they buy the core product too rarely—and while high add-on prices are the cause of the first problem they mitigate the second by funding a subsidized core product price. As such, controlling the maximum permitted add-on price need not benefit these naive consumers.

There are other situations where the presence of savvy consumers has no significant impact on the deals offered to the naive, and *vice versa*. For example, some consumers might not believe in the predictive power of horoscopes and ignore this market altogether, while others are willing to pay for this service. (We can remain neutral about who are the savvy here.) Unless there are strong scale economy effects (so that having large numbers of customers allows astrologers to operate more efficiently), there is no interaction between the two groups of consumers. More generally, many "scams" prey on the naive but have little impact on the savvy.

The phenomenon can also be seen in competitive insurance markets where some consumers are over-optimistic (or over-pessimistic) about the likelihood of the bad outcome.[39] Similarly, lenders may offer distinct contracts to entrepreneurs who are realistic or who are over-optimistic about their prospects.[40] Alternatively, naive consumers might be over-optimistic about how often they will go to an exercise gym. Such consumers may prefer a lump-sum membership, which is wrongly perceived to be "good value" by the optimistic consumer. A savvy consumer who accurately estimates her demand prefer a pay-per-visit contract, and neither type of consumer wishes to use the tariff aimed at the other type. Similar effects arise in situations where all consumers have self-control problems, and savvy consumers foresee this in advance while naive consumers do not.[41]

These various situations all share the same basic structure. The market is competitive, and so industry profit is zero. All consumers ultimately exhibit the same behaviour and impose the same costs on their supplier when faced with a given contract, and so the set of contracts which are consistent with zero profit is the same for a naive as for a savvy consumer. From this set, sellers in equilibrium choose the contract which is most attractive

[39] See Sandroni and Squintani (2007) for a model along these lines.
[40] See Landier and Thesmar (2009).
[41] See DellaVigna and Malmendier (2004) and Spiegler (2011, section 2.3) for further discussion.

ex ante to the target consumer, and so by construction, neither type is tempted by the contract aimed at the other group. The outcome is as if a consumer's savviness or naivete was known to sellers, and there is no interaction between the two groups.

3.3 Ripoff externalities

The final model is another model with "hidden costs" for naive consumers. Here, naive consumers mistakenly buy an add-on service which no one particularly wants or needs, while savvy can avoid the add-on costs. In this situation, the core price is subsidized with the profit generated by the fraction of naive consumers who end up paying for unwanted add-ons, and this benefits the savvy consumers who only pay the core price.[42]

Examples of the kind of add-on "service" I have in mind are as follows. Some airlines charge for carrying excess luggage, for checking-in luggage, or for checking in at the airport rather than online.[43] Savvy consumers are aware that these charges will be levied unless they take care in advance, while naive consumers will pay these charges if they turn up at the airport unprepared. Similarly, banks or credit card companies levy charges for unauthorized overdrafts or late payment. By being aware of their finances, savvy consumers can avoid these charges, while naive consumers might not be aware of the circumstances in which these charges can be levied.[44] Mobile phone contracts usually allow a specified number of calls per month, but if the subscriber makes more calls than this she pays an "overage" charge. Naive consumers who do not pay attention to their monthly usage or the possibility that they may need more than the monthly allowance may get caught out

[42] Gabaix and Laibson (2006) is perhaps the first and most prominent paper which discusses this phenomenon. Their approach differs slightly from that presented here. They suppose that a seller decides whether to advertise or to "shroud" its add-on price. When a seller advertises its price, this acts as an "eye-opener" and consumers realize they will have to pay the charge unless they take evasive action in advance. If sellers decide to shroud, they will choose monopolistic terms for the add-on. Savvy consumers anticipate this incentive, and take evasive action, while naive types do not. In many cases, an equilibrium exists in which all sellers shroud their add-on price, and naive consumers end up paying it. See Köszegi (2015, section 6) for a detailed survey of models of this sort.

[43] At the time of writing, *Ryanair* charges £70 to check in at the airport. See www.ryanair.com/en/fees for details (visited 21 May, 2014).

[44] Armstrong and Vickers (2012) discuss unauthorized overdraft fees in the UK. In the UK bank market only a minority of consumers pay such fees (which were an average of £23 per item in 2006), and these fees help fund the "free if in credit" model enjoyed by the majority of other consumers. (In 2006, about 30% of current account revenue came from these charges.) About 75% of account holders did not pay these fees, while 1.4 million customers paid more than £500 in such fees in 2006. The great majority of customers say they do not consider the level of these charges when choosing their bank, and few of those who paid these charges in 2006 anticipated beforehand having to pay them.

in a contract with high overage charges.[45]

Similar effects are seen in other scenarios. For instance, naive consumers might be susceptible to persuasion to buy a useless add-on. When a consumer buys a new television, a salesman may suggest she also buys an extended warranty to go with it. If the television is so reliable that the warranty actually has no value, the cost of the warranty is pure loss to the consumer and pure profit to the seller. Savvy consumers are immune to the salesman's patter and do not buy. Alternatively, practices such as "teaser" rates and roll-over contracts can be interpreted in a similar manner. Suppose that sellers supply a product over time, and the price for the first period's consumption is lower than for subsequent consumption. A savvy consumer might cancel her contract after one period (and perhaps enjoy another teaser rate from the next seller), while a naive consumer forgets to cancel or is unaware that her contract will automatically be rolled-over into the next period. This analysis is consistent with marketing tactics such as a bank offering a relatively high interest rate on a savings account for the first year, which drops off sharply thereafter, or a magazine offering a cheap trial period.[46]

To model these situations, consider the following stylized framework. Two or more sellers supply a product, the cost of which is C and the price of which is P. A fraction σ of consumers are savvy and pay only this price P. The remaining $1 - \sigma$ consumers are naive, and can be tricked into making an exogenous extra payment $R > 0$ to their chosen seller once they have purchased the product. This extra payment might be generated via small-print "traps" or worthless add-ons, which savvy consumers know how to avoid without cost. If sellers cannot distinguish the two kinds of consumers in advance, the equilibrium outcome in this market is for the product to be subsidized by the anticipated rents from the naive, so that

$$P = C - (1 - \sigma)R \ . \tag{10}$$

A savvy consumer pays only this bargain price, while a naive consumer pays the bargain price followed by the ripoff R, which comes to $C + \sigma R$ in total. Thus, both types of consumer pay more when σ is larger. A seller makes positive profit, equal to σR, from each non-savvy customer, but makes a loss $(1 - \sigma)R$ from each savvy customer, and on

[45] See Grubb (2015a) for analysis of this form of bill shock, and the possible interventions to overcome the problem.

[46] Even the consumer rights body in the UK, *Which?*, employs this tactic. One can sign up for one month's service for just £1, which is automatically rolled-over for £10.75 each month until cancelled. See www.which.co.uk/signup for further details (visited 21 May 2014).

average it just breaks even.

As before, a consumer has idiosyncratic valuation X for the product, $Q(P)$ is the proportion of consumers with $X \geq P$ and $S(P) \equiv \int_P Q$ measures consumer surplus from the product (without ripoffs) when price is P. The average surplus of a savvy consumer in this market is then $V_S(\sigma) = S(C - (1 - \sigma)R)$, which decreases with σ. The (true) surplus of a naive consumer is

$$V_N(\sigma) = S(C - (1 - \sigma)R) - RQ(C - (1 - \sigma)R)$$

since a naive consumer buys just as often as a savvy type, but ends up paying an extra R if she does buy. (Here, it is possible that V_N is negative.) It is ambiguous whether or not V_N decreases with σ. However, if R is not too large or demand Q is not too elastic, V_N will, like V_S, decrease with σ so that a ripoff externality is present. Aggregate consumer surplus, $V(\sigma) = \sigma V_S(\sigma) + (1 - \sigma)V_N(\sigma)$, which equals total welfare in this competitive market, unambiguously *increases* with σ due to the larger number of consumers who enjoy the higher surplus V_S.

Aggregate consumer surplus $V(\sigma)$ always falls with R, and so there is scope for welfare-improving regulation which constrains the size of the ripoff R. However, the impact of such regulation on the two groups of consumers differs: a savvy type benefits from a seller's ability to rip off the naive and so would like R to remain large, while a naive consumer's surplus V_N decreases with R. As such, the two groups have opposing interests towards regulation to limit ripoffs, and savvy types—who might (as in footnote 44) be in the majority—have an incentive to lobby against this welfare-enhancing regulation.[47]

Finally, it may be possible for sellers to observe, or learn about, a consumer's savviness, and set prices accordingly. For instance, banks or credit card companies may have information about a consumer's propensity to pay penalty fees. In this case, if banks can set personalized prices to different consumers, they would charge more for their core service to savvy consumers who avoid extra fees than to those consumers who are known to pay

[47] If R is large enough, the price in (10) is negative. If a negative price is not feasible, the outcome is then that the product is offered for free, and the firm's costs are covered entirely by exploiting the naive. (For instance, in the UK a bank account is typically free.) Sellers are then less able to dissipate profits, and profits may be positive even in a competitive market. In these cases, sellers as well as savvy consumers have an incentive to lobby against constraints on ripoffs. However, in the credit card and banking contexts, it may be possible for sellers to in effect set a negative price for their product. For example, "cashback" contracts (where a consumer is paid by her credit card for each transaction) or current accounts which pay interest to customers have this flavour.

them more frequently. This possibility reduces the linkage between the two groups, and in the limit externalities are eliminated as in sections 2.2 and 3.2.

4 Conclusions

This paper has explored how the balance of "savvy" and "non-savvy" consumers in a market affects the deals which sellers offer their customers. I discussed three ways in which the two groups might interact: the case of search externalities, where savvy consumers help all consumers obtain a good deal; the case of ripoff externalities, where non-savvy consumers facilitate good deals for all consumers, and the case without interactions between the two groups.

I restricted attention to two broad kinds of market: those which exhibit price dispersion (section 2) and those involving forms of add-on pricing (section 3). The classical model of price dispersion in section 2.1 provided an instance of the search externality. There, savvy consumers shopped around for the lowest price and their demand was more elastic. As such, savvy consumers induced lower prices in equilibrium, so that industry profit fell with σ, the fraction of savvy types, while overall welfare increased. Because profit falls when more consumers are savvy, sellers have an incentive to try to confuse consumers in the way they present their offers, and they may welcome regulation which reduces the incentive for consumers to become savvy.

At the other extreme was the model of "bill shock" in section 3.3, where non-savvy consumers could be tricked into making extra payments. Competing sellers set the price for their product in anticipation that a consumer might be a naive type and generate extra revenue. As such, the equilibrium price for the product was subsidized, with a greater subsidy when the fraction of non-savvy types was larger. In many cases, both types of consumers benefit from the presence of the naive, and a ripoff externality was present. Overall welfare decreases with the ability to exploit naive consumers. Regulation to constrain ripoffs may therefore be efficient, but may be resisted by savvy types who prey on the naive consumers when they stumble into small-print traps.

The framework followed in this paper can be applied to several other kinds of markets, and to conclude I list four further applications. First, a likely example of a market in which non-savvy consumers might help all consumers is Akerlof (1970)'s lemons market, where savvy consumers who understand adverse selection can cause the market to shut down. Strategically naive consumers, however, who mistakenly believe the pool of products offered for sale is unaffected by the selling price—and who may therefore pay more than the product is really worth to them—can allow the market to function.[48]

Second, consider a simple model of hold-up along the lines of Diamond (1971)'s paradox. Suppose a single firm supplies an indivisible product and all consumers must incur a sunk cost to travel to the seller to discover its price. If all consumers knew their idiosyncratic value for the item in advance, and anticipate the firm's incentive to set its price, this market would break down entirely. (The firm always has an incentive to raise its price a little above any price anticipated by consumers, since a consumer travels to the seller if her value is greater than the price plus the travel cost, and so no equilibrium exists where some consumers travel to the firm.) However, if some non-savvy consumers do not know their value for the item until they travel to the seller, this can enable the market to open, to the benefit of all consumers.[49]

Consider next examples of markets which plausibly exhibit a search externality. First, savvy consumers may be less willing to pay a brand premium when the underlying product quality does not justify this. For instance, Bronnenberg, Dubé, Gentzkow, and Shapiro (2015) document how doctors are less likely than the average consumer to buy brand-name headache remedies, which may be much more expensive than a pharmacy's generic equivalent. (They find that only 9% of doctors buy a branded product over the generic, chemically identical, substitutes, while 26% of other consumers do so.) In a calibrated model, the authors suggest that brand-name prices would fall if more consumers behaved like doctors, and this would benefit all consumers in the market.

Finally, consider the durable goods problem of Coase (1972), where a single firm sells its product over time to forward-looking consumers with heterogeneous tastes for its product. The firm cannot commit to its future prices, and after high-value consumers have purchased, the firm has an incentive to reduce its price to sell to remaining lower-value consumers. This model can be viewed as an oligopoly market—where the firm "competes with itself" over time—with inter-temporal price dispersion. The model can be extended so that some consumers are naive, in that they do not foresee the seller's incentive to reduce its price over time. As such, these naive consumers buy myopically, as soon as the price

[48] See Spiegler (2011, section 8.3) and the references listed there for further discussion of markets when consumers have limited understanding of adverse selection.

[49] See Anderson and Renault (2006) for a model which makes use of this insight.

falls below their valuation for the item. These naive consumers are like the inert shoppers in Varian's model (who can be interpreted as mistakenly believing that all firms offer the same price). The presence of these consumers tends to relax intra-firm competition in Coase's model, just as they relax inter-firm competition in Varian's model, and so this market is likely to involve a search externality.

References

Akerlof, G. (1970): "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84(3), 488–500.

Anderson, S., and R. Renault (2000): "Consumer Information and Firm Pricing: Negative Externalities from Improved Information," *International Economic Review*, 41(3), 721–742.

——— (2006): "Advertising Content," *American Economic Review*, 96(1), 93–113.

Armstrong, M. (2008): "Interactions between Competition and Consumer Policy," *Competition Policy International*, 4(1), 97–147.

Armstrong, M., and J. Vickers (2012): "Consumer Protection and Contingent Charges," *Journal of Economic Literature*, 50(2), 477–493.

Armstrong, M., J. Vickers, and J. Zhou (2009): "Consumer Protection and the Incentive to Become Informed," *Journal of the European Economic Association*, 7(2-3), 399–410.

Armstrong, M., and J. Zhou (2011): "Paying for Prominence," *Economic Journal*, 121(556), F368–395.

Baye, M., and J. Morgan (2002): "Information Gatekeepers and Price Discrimination on the Internet," *Economics Letters*, 76(1), 47–51.

Bronnenberg, B., J.-P. Dubé, M. Gentzkow, and J. Shapiro (2015): "Do Pharmacists Buy Bayer? Informed Shoppers and the Brand Premium," mimeo.

Brown, J., and A. Goolsbee (2002): "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry," *Journal of Political Economy*, 110(3), 481–507.

Burdett, K., and K. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51(4), 955–969.

Chioveanu, I., and J. Zhou (2013): "Price Competition with Consumer Confusion," *Management Science*, 59(11), 2450–2469.

Clerides, S., and P. Courty (2013): "Sales, Quantity Surcharge, and Consumer Inattention," mimeo.

Coase, R. (1972): "Durability and Monopoly," *Journal of Law and Economics*, 15(1), 143–149.

Crosetto, P., and A. Gaudeul (2014): "Choosing Whether to Compete: Price and Format Competition with Consumer Confusion," Jena Economic Research Papers 2014-026.

DellaVigna, S., and U. Malmendier (2004): "Contract Design and Self-Control: Theory and Evidence," *Quarterly Journal of Economics*, 119, 353–402.

Diamond, P. (1971): "A Model of Price Adjustment," *Journal of Economic Theory*, 3(2), 156–168.

Ellison, G. (2005): "A Model of Add-on Pricing," *Quarterly Journal of Economics*, 120(2), 585–637.

Fershtman, C., and A. Fishman (1994): "The Perverse Effects of Wage and Price Controls in Search Markets," *European Economic Review*, 38(5), 1099–1112.

Gabaix, X., and D. Laibson (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, 121(2), 505–540.

Grossman, S., and J. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 79(3), 393–408.

Grubb, M. (2015a): "Consumer Inattention and Bill-Shock Regulation," *Review of Economic Studies*, 82(1), 219–257.

——— (2015b): "Overconfident Consumers in the Marketplace," mimeo.

Knittel, C., and V. Stango (2003): "Price Ceilings as Focal Points for Tacit Collusion: Evidence from Credit Cards," *American Economic Review*, 93(5), 1703–1729.

Koszegi, B. (2015): "Behavioral Contract Theory," *Journal of Economic Literature*, 52(4), 1075–1118.

Landier, A., and D. Thesmar (2009): "Financial Contracting with Optimistic Entrepreneurs," *Review of Financial Studies*, 22(1), 117–150.

McAfee, R. P., and M. Schwartz (1994): "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity and Uniformity," *American Economic Review*, 84(1), 210–230.

Miravete, E. (2013): "Competition and the Use of Foggy Pricing," *American Economic Journal: Microeconomics*, 5(1), 194–216.

Morgan, J., H. Orzen, and M. Sefton (2006): "An Experimental Study of Price Dispersion," *Games and Economic Behavior*, 54(1), 134–158.

Petrikaite, V. (2014): "Collusion with Costly Consumer Search," mimeo, University of Groningen.

Piccione, M., and R. Spiegler (2012): "Price Competition Under Limited Comparability," *Quarterly Journal of Economics*, 127(1), 97–135.

Salop, S., and J. Stiglitz (1977): "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies*, 44(3), 493–510.

Sandroni, A., and F. Squintani (2007): "Overconfidence, Insurance and Paternalism," *American Economic Review*, 97(5), 1994–2004.

Schultz, C. (2005): "Transparency on the Consumer Side and Tacit Collusion," *European Economic Review*, 49(2), 279–297.

Scitovsky, T. (1950): "Ignorance as a Source of Oligopoly Power," *American Economic Review*, 40(2), 48–53.

Shapiro, C. (1995): "Aftermarkets and Consumer Welfare: Making Sense of Kodak," *Antitrust Law Journal*, 62(2), 484–511.

Spiegler, R. (2011): *Bounded Rationality and Industrial Organization*. Oxford University Press, Oxford, UK.

Varian, H. (1980): "A Model of Sales," *American Economic Review*, 70(4), 651–659.

Wilson, C. (2010): "Ordered Search and Equilibrium Obfuscation," *International Journal of Industrial Organization*, 28(5), 496–506.

Technical Appendix

Details for the model presented in section 2.3: Depending on the proportion of savvy consumers, there are two kinds of equilibria to consider: those where sellers only ever serve the savvy consumers, and those where sellers sometimes serve non-savvy consumers.[50]

First, we look for an equilibrium with price support $[p_{\min}, v]$, where $p_{\min} > \alpha v$ so that non-savvy consumers never participate in the market. The equilibrium profit for a seller is then $\sigma\alpha v(1-\alpha)$. If a seller deviated and instead set the lower price $p = \alpha v$ it would sell to $1-\sigma+\sigma\alpha$ consumers, and so obtain profit $\alpha v[1-\sigma+\alpha\sigma]$. For this deviation to be unprofitable we require $\sigma\alpha v(1-\alpha) \geq \alpha v[1-\sigma+\alpha\sigma]$, or

$$\sigma(1-\alpha) \geq \tfrac{1}{2} \; . \tag{11}$$

If $\alpha > \frac{1}{2}$ this inequality is never satisfied, while if $\alpha \leq \frac{1}{2}$ it is satisfied if and only if σ is large enough. The CDF for a seller's mixed strategy for price, F, satisfies

$$\sigma\alpha\left[1-\alpha+\alpha(1-F(p))\right]p \equiv \sigma\alpha v(1-\alpha) \; , \tag{12}$$

so that F does not depend on σ. The minimum price in the support is therefore given by

$$p_{\min} = (1-\alpha)v \; ,$$

[50] There is no equilibrium where sellers *always* serve the uninformed, and so set price no higher than αv. To see this, suppose such an equilibrium did exist. If so, the maximum price chosen in equilibrium is αv, and equilibrium profit for a seller is $\sigma\alpha^2 v(1-\alpha)$, which is its profit from a savvy consumer who has a good match with it but not with its rival. If instead this seller set price $p = v$ it would sell to the same consumers, but make more profit. We deduce there is no such equilibrium.

which from (11) is indeed greater than αv as required.

Next, look for an equilibrium where the minimum price satisfies $p_{\min} < \alpha v$, so that non-savvy consumers are sometimes served. In this case, the price support will consist of two disjoint segments, with a gap in between, so that the price support takes the form $[p_{\min}, \alpha v] \cup [p^*, v]$ where $p^* > \alpha v$.[51] Again, a seller obtains equilibrium profit $\sigma\alpha v(1-\alpha)$, and so $p_{\min}$ satisfies

$$[1-\sigma+\sigma\alpha]\, p_{\min} = \sigma\alpha v(1-\alpha)$$

which indeed implies that $p_{\min} \leq \alpha v$ when (11) does not hold. Let $F(p)$ denote the CDF for each seller's mixed strategy for choosing price. In the higher segment $[p^*, v]$ we again have F satisfying (12). In the lower segment $[p_{\min}, \alpha v]$ we have

$$[\sigma\alpha(1-\alpha) + (\sigma\alpha^2 + 1 - \sigma)(1 - F(p))]p = \sigma\alpha v(1-\alpha) \ . \tag{13}$$

For a seller to be indifferent between choosing price $p = p^*$ and $p = \alpha v$, we require

$$[\sigma\alpha(1-\alpha) + (\sigma\alpha^2 + 1 - \sigma)(1 - F(p^*))]\alpha v = \sigma\alpha v(1-\alpha) \ ,$$

so that the probability a seller caters only to the savvy, $1 - F(p^*)$, satisfies

$$1 - F(p^*) = \frac{\sigma(1-\alpha)^2}{1-\sigma(1-\alpha^2)} \ .$$

(Note that $1 - F(p^*) = 1$ when (11) just binds.) It follows from (12) that the price p^* satisfies

$$\left[1-\alpha+\alpha\frac{\sigma(1-\alpha)^2}{1-\sigma(1-\alpha^2)}\right] p^* = v(1-\alpha)$$

or

$$p^* = \frac{1-\sigma+\alpha^2\sigma}{1-\sigma+\alpha\sigma} v$$

which is always greater than αv. Note that p^* decreases with σ. This completes the derivation of the equilibrium. Regardless of whether condition (11) holds, industry profit is given by expression (4) in the text.

In this equilibrium, a seller's price weakly increases with σ in the sense of first-order stochastic dominance. For instance, if (11) holds, F does not depend on σ, while if (11) does not hold one can check that F in (13) decreases with σ. Because of this, a consumer

[51] A seller's demand jumps at the point $p = \alpha v$, since for a price below αv a seller serves all non-savvy consumers when its rival sets price above αv. The gap in the support is needed so that a seller can be indifferent between all prices in the support.

of either type is better off when σ falls. Consider the particular case with $\alpha = \frac{1}{2}$ and $v = 2$ in more detail. Here, condition (11) never holds, and the non-savvy always have a chance to purchase. The following figures depict the equilibrium price density when $\sigma = \frac{1}{3}$ and $\sigma = \frac{2}{3}$ respectively, and shows how increasing σ induces sellers to set higher prices.

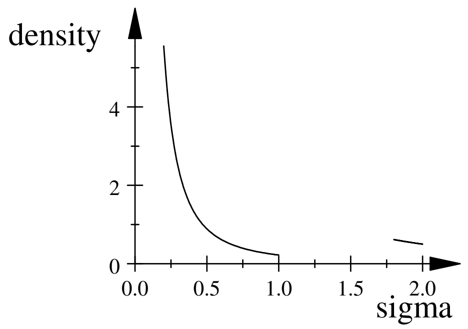

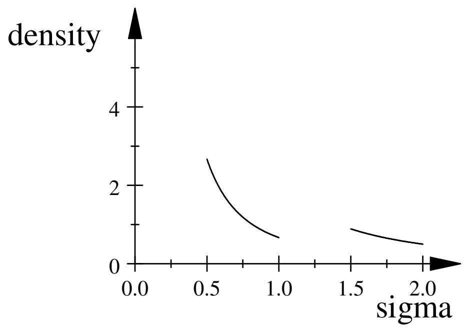

Figure 4: The impact of increasing σ on the density for prices

The expected surplus of a non-savvy consumer is

$$V_N(\sigma) = \int_{p_{\min}}^{\alpha v} (\alpha v - p) dF_{\min}(p) = \int_{p_{\min}}^{\alpha v} F_{\min}(p) dp = \int_{p_{\min}}^{\alpha v} (1 - (1 - F(p))^2) dp \ .$$

Here, $F_{\min}(p)$ is the CDF for the minimum of the two prices, which equals $1 - (1 - F(p))^2$. In the example where $v = 2$ and $\alpha = \frac{1}{2}$, we have $p_{\min} = \frac{\sigma}{2-\sigma}$ and (13) implies that in the range $[p_{\min}, 1]$ we have $1 - F(p) = \frac{\sigma}{4-3\sigma}(\frac{2}{p} - 1)$, and the resulting $V_N(\sigma)$ above is depicted as the lower solid curve on Figure 3 in the text.

If condition (11) holds, total welfare is $W(\sigma) = \sigma v(1 - (1 - \alpha)^2)$, while if (11) does not hold the non-savvy sometimes buy and welfare is then $W(\sigma) = (1-\sigma)\alpha v\,(1 - (1 - F(p^*))^2) + \sigma v(1 - (1 - \alpha)^2)$. In the example where $v = 2$ and $\alpha = \frac{1}{2}$, the latter case applies and $1 - F(p^*) = \frac{\sigma}{4-3\sigma}$. Aggregate consumer surplus, $V(\sigma)$, is equal to $W(\sigma) - \Pi(\sigma)$, and in Figure 3 this is plotted as the dashed curve. Finally, savvy surplus V_S can be obtained from the formula $\sigma V_S(\sigma) + (1 - \sigma)V_N(\sigma) = V(\sigma)$, and is plotted as the higher solid curve on Figure 3.

Details for the model presented in section 3.1: Given σ, suppose that (P^*, p^*) is the equilibrium price pair in this market, and write $U = s(p^*) - P^*$ for consumer surplus gross of transport costs in this equilibrium. Suppose one seller considers choosing the price pair

(P, p), while its rival plays the equilibrium strategy (P^*, p^*). Savvy consumers see both of this seller's prices, and so the fraction who buy from it is

$$\frac{1}{2} + \frac{s(p) - P - U}{2t} .$$

Non-savvy consumers do not observe p, and form rational expectations about this price given P. Suppose they anticipate add-on price $\hat{p}$, in which case the fraction of non-savvy who buy from the seller is

$$\frac{1}{2} + \frac{s(\hat{p}) - P - U}{2t} .$$

Similarly to expression (5), the seller's total profit from the two groups is therefore

$$\left[\sigma\left(\frac{1}{2} + \frac{s(p) - P - U}{2t}\right) + (1-\sigma)\left(\frac{1}{2} + \frac{s(\hat{p}) - P - U}{2t}\right)\right] \times [P - C + \pi(p)] . \qquad (14)$$

Given P and $\hat{p}$, the seller chooses its add-on price p to maximize this profit. (The seller will never choose an add-on price above the monopoly price p^M in problem (14), regardless of consumer beliefs $\hat{p}$.) For the non-savvy consumers to have rational expectations about the add-on price, we require that choosing $p = \hat{p}$ maximizes profit in (14). This implies that the equilibrium add-on price given P satisfies the first-order condition

$$(t + s(p) - P - U)\pi'(p) - \sigma q(p)(P - C + \pi(p)) = 0 . \qquad (15)$$

Given equilibrium surplus U, expression (15) characterizes the relationship (even off the equilibrium path) between a seller's core price and its rationally anticipated add-on price.

Note that if $\sigma = 0$, the equilibrium add-on price satisfies $\pi'(p) = 0$, so that $p = p^M$ regardless of the core price P and the equilibrium add-on price is $p^* = p^M$. At the other extreme, if $\sigma = 1$ the outcome as discussed in the text involves $p^* = c$, $P^* = C + t$ and $U = s(p^*) - P^*$, which satisfies (15).

Expression (15) can be written as

$$(1 - \eta(p))(t + s(p) - P - U) = \sigma(P - C + \pi(p)) , \qquad (16)$$

where $\eta(p) \equiv -pq'(p)/q(p)$ is the elasticity of add-on demand (which is increasing given that q is logconcave). Since η is increasing, the left-hand side of (16) is decreasing with p (at least in the range $c \le p \le p^M$), while the right-hand side is increasing. Thus, there is a unique price p which solves (16). Since the left-hand side of (16) decreases with P, while the right-hand side is increasing with P, it follows that the equilibrium add-on price

p is a decreasing function of P. Intuitively, when the seller undercuts the equilibrium core price, this reduces its incentive to gain market share from the savvy consumers with a low add-on price. This is the crucial insight which reveals why the market is less competitive when sellers face a mixed population of consumers.

If $\sigma > 0$, from (15) we can write the core price P which implements a given add-on price p as

$$P(p) = \frac{\pi'(p)(t + s(p) - U) + \sigma q(p)(C - \pi(p))}{\pi'(p) + \sigma q(p)} \ . \tag{17}$$

The seller then chooses its add-on price p and associated core price $P(p)$ to maximize its profit

$$\left[\frac{1}{2} + \frac{s(p) - P(p) - U}{2t}\right] \times [P(p) - C + \pi(p)] \ . \tag{18}$$

The remaining condition is to ensure that U in (17)–(18) is the equilibrium level of consumer surplus, i.e., that $U = s(p^*) - P^*$. In sum, (P^*, p^*) constitutes a symmetric equilibrium if:

(i) p^*maximizes (18), where $P(p)$ is defined in (17);
(ii) $P^* = P(p^*)$, and
(iii) $U = s(p^*) - P^*$.

It is hard to make further progress at this level of generality, and so I illustrate the analysis with the example where $q(p) = 1 - p$ and $C = c = 0$, so that $s(p) = \frac{1}{2}(1-p)^2$, $\pi(p) = p(1-p)$ and $p^M = \frac{1}{2}$. Expression (17) then implies

$$P(p) = \frac{(1-2p)(t + \frac{1}{2}(1-p)^2 - U) - \sigma p(1-p)^2}{1 - 2p + \sigma(1-p)} \ .$$

Substituting this into profit (18) shows that the seller's profit with add-on price p is

$$\frac{1}{8t}\sigma(1-p)(1-2p)\left(\frac{1-p^2-2U+2t}{1-2p+\sigma(1-p)}\right)^2 \ ,$$

and the (relevant) first-order condition for maximizing this expression is that p satisfies

$$((1-p)(1-2p))' \frac{1-p^2-2U+2t}{1-2p+\sigma(1-p)} + 2(1-p)(1-2p)\left(\frac{1-p^2-2U+2t}{1-2p+\sigma(1-p)}\right)' = 0$$

or

$$(4p-3)(1-p^2-2U+2t) + 2(1-p)(1-2p)\left(-2p + (2+\sigma)(\frac{1-p^2-2U+2t}{1-2p+\sigma(1-p)})\right) = 0 \ . \tag{19}$$

Given that $U = \frac{1}{2}(1-p)^2 - P$, expression (17) implies that

$$U = \frac{\sigma(1+p)(1-p)^2 - 2t(1-2p)}{2\sigma(1-p)} \ , \tag{20}$$

and substituting this value of U into (19) reveals that the equilibrium add-on price p^* given σ satisfies

$$t(1-2p^*) - \sigma(1-p^*)\left(t + 2p^*(1-p^*)(1-2p^*)\right) = 0 \ . \tag{21}$$

When $\sigma = 0$, this expression has solution $p^* = \frac{1}{2}$, while when $\sigma = 1$ it has solution $p^* = 0$. From (21), we can write the σ which implements a given add-on price $p^* \in [0, \frac{1}{2}]$ as

$$\sigma = \frac{t(1-2p^*)}{(1-p^*)\left(t + 2p^*(1-p^*)(1-2p^*)\right)} \ . \tag{22}$$

Unless t is very small indeed (below 0.01 or so), the above expression decreases with p^* in the range $p^* \in [0, \frac{1}{2}]$, and so there is a one-to-one decreasing relationship between σ and the equilibrium add-on price p^*. When σ satisfies (22), equilibrium consumer surplus with add-on price p^*, which is U in (20), simplifies to

$$U = \tfrac{1}{2}(1 - 8[p^*]^3 + 11[p^*]^2 - 4p^*) - t \ .$$

This is "U-shaped" in p^* in the relevant range $p \in [0, \frac{1}{2}]$, and hence U-shaped as a function of σ. This market has a ripoff externality when the initial σ is small, while for large σ the market has a search externality. (However, consumer surplus when $\sigma = 1$, i.e., when $p = 0$, is above that when $\sigma = 0$, i.e., when $p = \frac{1}{2}$.) Likewise, industry profit when the equilibrium add-on price is p^*, $P(p^*) + \pi(p^*)$, in this example is

$$2p^*(1-p^*)(1-2p^*) + t \ ,$$

which is hump-shaped in p^*. Thus, industry profit is maximized with a mixed population of consumers, rather than when all consumers are savvy or all are non-savvy. It is hard to invert expression (22) to obtain the add-on price p^* as an explicit function of σ. It is straightforward to obtain numerical results, however, and Table 1 in the text reports outcomes for a range of σ when $t = 1$.