

MPRA

Munich Personal RePEc Archive

Search and Ripoff Externalities

Mark Armstrong

Department of Economics, University of Oxford

July 2014

Online at <https://mpra.ub.uni-muenchen.de/57958/>

MPRA Paper No. 57958, posted 17 August 2014 04:52 UTC

Search and Ripoff Externalities

Mark Armstrong*

July 2014

Abstract

This paper surveys models of markets in which some consumers are “savvy” while others are not. We discuss when the presence of savvy consumers improves the deals available to non-savvy consumers in the market (the case of search externalities), and when the non-savvy fund generous deals for savvy consumers (ripoff externalities). We also discuss when the two groups of consumers have aligned or divergent views about market interventions. The analysis covers two overlapping families of models: those which examine markets with price/quality dispersion, and those which exhibit forms of consumer hold-up.

Keywords: Consumer protection, consumer search, price dispersion, hold-up, add-on pricing.

1 Introduction

This paper examines a number of 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 that consumer policies which protect vulnerable consumers are only needed 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 better informed about the deals available in the market ensure that less informed consumers also obtain reasonable outcomes. Recent work, however, has examined situations where savvy consumers benefit from the presence of non-savvy types. In such markets, a “ripoff externality” is present, and vulnerable consumers may need protection even when they are small in number.

*Department of Economics and All Souls College, University of Oxford. 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, Rani Spiegler, John Vickers, and Jidong Zhou.

This paper discusses three principal issues. First, what kinds of inter-consumer externalities are present? 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 intervention benefit both consumer groups and which policies benefit one group at the expense of the other? Third and finally, what determines the extent of savviness in the consumer population?

For our purposes, there are two broad notions of savviness to consider. First, a consumer might be *well informed* about the prices and/or product qualities available in the market. For instance, a savvy consumer shopping for wine is able to determine the likely quality of the wine inside by looking at the label. Alternatively, a consumer looking for a new television may know the range of available prices (e.g., because she is online), or knows how much she is willing to pay for a product before travelling to the seller. Second, 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 (i.e., they are non-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 how she will behave.

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. For instance, in a market where in fact there is price dispersion, but naive consumers think that all sellers offer the same price, a naive consumer might choose not to incur search costs to become informed about the prices in the market.

A useful framework for discussing the 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, we suppose that there are no systematic differences in *tastes* for the product in question across the two groups of consumers. Except for section 2.2, we take the extent of savviness, σ , to be exogenous and out of the control of consumers and firms.

Let $V_S(\sigma)$ and $V_N(\sigma)$ denote the expected net surplus enjoyed in equilibrium by an

¹See also Armstrong (2008, section III.C) and Armstrong and Vickers (2012, section 3).

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 better deals than their non-savvy counterparts.² 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 . 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 all profits to be zero. Finally, let $W(\sigma) \equiv V(\sigma) + \Pi(\sigma)$ denote total welfare.

There are (at least) three cases of interest:

Search externalities: When savvy consumers exert a positive externality on the non-savvy—that is, when $V_N(\sigma)$ increases with σ —we 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, and when there are more consumers aware of all the available deals this makes suppliers offer good deals, which in turn are available to more inert buyers.⁴ As we will see, within this class of markets, V_S usually also increases with σ , while overall welfare W can increase or decrease with σ and profits Π might increase, decrease or be “hump-shaped” in σ , depending on the context.

Ripoff externalities: When savvy consumers benefit from the presence of the non-savvy—

²However, there are situations in which replacing a population of savvy buyers with a population of non-savvy buyers will make buyers better off. For instance, this is the case in one of the hold-up scenarios discussed in section 3.1. There are also cases where the two kinds of consumer obtain exactly the *same* surplus; for example, this is often the case when a monopolist offers a single product at a single price and so all consumers obtain the same deal.

³For instance, in section 2.3 the surplus enjoyed by savvy consumers can be a non-monotonic function of σ , although non-savvy surplus increases with σ . Likewise, in the model of “bill shock” in section 3.2, it is possible that V_N increases with σ while V_S decreases with σ .

⁴At the time of writing this, the front-page headline of the UK’s *Daily Telegraph* on 9 July 2014 was “Savvy shoppers force down prices”.

when $V_S(\sigma)$ decreases with σ —we 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. Another example of such a market (not discussed further in this paper) 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 open.⁵ In markets with ripoff externalities, it is possible that aggregate consumer surplus V rises with σ , even if both V_S and V_N fall with σ , if the gap $V_S - V_N$ is large.

No interactions between consumers: On the knife edge between these two 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 typically involve biased beliefs on the part of naive consumers. If present, 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, biased consumers might regret the deal they chose. (A lucky charm which is sold 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 naive consumers.)

The plan for the rest of this paper is as follows. Oligopoly models which generate price or quality dispersion are examined in section 2, and we will see that the search externality tends to operate in such markets, so that savvy types confer a benefit to the non-savvy (and usually to other savvy types too). Models with various forms of “hold-up” are presented in section 3, including situations with both an indivisible good and with a more complex product involving add-on services.⁶ In these markets a richer set of outcomes are possible, and seemingly minor variants of the add-on price problem generate each of the three situations—search externalities, ripoff externalities, and no interaction—listed above. We end the paper with some concluding comments.

⁵See Spiegler (2011, section 8.3) and the references listed there for further discussion of markets when consumers have limited understanding of adverse selection.

⁶There is some overlap in the two classes of model. The model of quality dispersion in section 2.4 could logically fit under either heading, and one of the add-on models in section 3.2 involves price dispersion.

2 Price and Quality Dispersion

2.1 A model of price dispersion

In a market for an indivisible good of known quality, it is intuitive that when some consumers are aware of available prices and buy from the cheapest supplier, those who shop less diligently are partially protected.

To illustrate this, consider Varian (1980)'s classical model of price dispersion.⁷ Here, n identical firms supply a homogeneous product with unit cost c . In general, consumers 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 we assume is quasi-concave in p , and p^M for the price which maximizes this profit. An exogenous fraction σ of consumers (independent of the valuation v) are savvy, in the sense that they buy from the cheapest supplier, while other $1 - \sigma$ consumers buy from a random supplier so long as that supplier's price is below their v .⁸

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 supplier has an incentive to set price below the monopoly price p^M , and the outcome is as if a single firm 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 firm choosing its price according to a cumulative distribution function (CDF) $F(p)$, which satisfies

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

⁷See Salop and Stiglitz (1977) for closely related analysis.

⁸This behaviour could be justified if each consumer's cost of search is very convex, in the sense that a consumer can visit one supplier for free but finds it too costly ever to visit a second supplier. A fraction σ are informed of each firm'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 think that competition operates so that the "law of one price" operates and all sellers offer the same price.

Here, a firm 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 firm will always sell to its share of the $1 - \sigma$ inert consumers (who have $v \geq p$). As such, a firm's demand from the uninformed consumers is less elastic than demand from the informed. 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.⁹ For a firm to be willing to play the mixed strategy $F(\cdot)$, the firm must be indifferent between all prices in the support of $F(\cdot)$.

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 makes a seller's demand more elastic. Because each seller's price distribution is shifted downwards (in the sense of first-order stochastic dominance) when σ rises, both the savvy consumers (who pay the minimum price from n draws) and the inert consumers (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. From (1), industry profit is $\Pi(\sigma) = (1 - \sigma)\pi(p^M)$, which decreases with σ . Total welfare W at least weakly increases with σ since lower prices stimulate demand.




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

⁹However, it is not an equilibrium for sellers to choose the monopoly price $p = p^M$ for sure, since a seller could slightly undercut this price and thereby serve all the savvy consumers.

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 in the market, is non-monotonic in σ . (As discussed, there is no price dispersion when $\sigma = 0$ 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.¹⁰

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 suppliers is larger, however, one can show that the expected price paid by savvy consumers falls while the expected price paid by the inert shoppers rises, so the two curves on Figure 1 move 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 inert, $(1 - \sigma)/n$, 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 suppliers has contrasting effects on the informed and the uninformed consumers, with only the savvy benefitting from "more competition" of this form.¹¹

Extension to this benchmark model:

A modification to the above model allows suppliers to charge distinct prices to savvy and inert 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 supplier 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 are charged 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 of either group.¹²

¹⁰Brown 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, then decreases as their use becomes more widespread.

¹¹See 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.

¹²Baye and Morgan (2002) consider a model in which sellers must pay to list on a price comparison

A second variant of Varian’s model extends the analysis to a dynamic setting, and examines the impact of consumer savviness on the sustainability of tacit collusion in this market.¹³ Suppose the industry attempts to collude at the monopoly price p^M with the use of a trigger strategy. If a firm deviates by undercutting p^M , suppose 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-firm profits in each period given by the right-hand side of (1). If a firm does undercut the collusive price, this lower price is observed only by the σ savvy consumers. As a result, when δ is the discount factor, collusion at the monopoly price can be sustained if

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

which reduces to the usual condition

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

In this market, increasing σ has two contrasting effects. When σ is large there is fierce competition without collusion, and so the punishment profits are small. On the other hand, when σ is large, the number of consumers who respond to a price cut is large, and so short-run gains from deviating are large. These two effects precisely cancel out, and the ability to collude is unaffected by the number of savvy consumers.

A final variant considers a situation in which, instead of purchasing from a *random* seller, a sales intermediary (or “salesman” for brevity) steers the inert consumers towards a supplier of his choice, if given incentive by that supplier to do so. These naive consumers follow sales advice, without understanding that the advice might be biased by financial inducements from sellers.¹⁴

website, and can charge different prices on this website and when they sell direct to consumers. They find that sellers choose whether to list according to a mixed strategy and choose their price on the comparison website according to a mixed strategy, and obtain positive profits there generated by the possibility they are the sole listing seller. The price on the comparison website is lower than its price on its own platform.

¹³See 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 can become informed about prices and valuations by incurring a search cost. She finds that an increase in this search cost—i.e., a reduction in market transparency—makes collusion easier to achieve.

¹⁴Inderst and Ottaviani (2012, page 502) report how a majority of people who had received financial advice believed that advice to be independent, and only a minority believed that commissions were being paid to their advisor.

In more detail Armstrong and Zhou (2011, section 1) suppose that a number of symmetric suppliers costlessly supply a homogenous product which all consumers value at v . This product is only available via a consultation with a salesman. An exogenous fraction σ of savvy consumers are immune to the salesman's patter, costlessly observe the full list of retail prices, and buy from the cheapest supplier. The remaining fraction $1 - \sigma$ of consumers are susceptible to the marketing efforts of the salesman and follow his recommendation. Suppose that a supplier chooses its retail price, p , and commission rate, b , simultaneously (and simultaneously with its rivals). In this setting a salesman will promote the highest-commission product (regardless of how retail prices compare).

When $\sigma = 1$, so that all consumers are savvy, there is no point in a seller spending resources to influence a salesman, and the result is Bertrand price competition, and suppliers and salesmen obtain no profits. When $\sigma = 0$, the salesman determines demand entirely, and so suppliers compete to offer the highest commission. The result is that both the retail price and the commission payment is driven up to v , so that suppliers obtain zero profit but salesmen extract the entire social surplus. In either case, there is no price dispersion.

When $0 < \sigma < 1$ sellers choose their retail prices and commission payments randomly. In equilibrium, there is an increasing relationship between a firm's choice of b and p . This is because a higher price p makes it more worthwhile for a seller to pay the salesman to steer the uninformed consumers towards its product. Since high commissions are associated with high retail prices, there is "mis-selling", and a salesman promotes the more expensive product due to the higher commission he receives. The expected outlay for a non-savvy consumer is the expected value of the *highest* of the retail prices in the market, rather than the the expected value of a random price in the market as in Varian's model.

In the case with two suppliers, Armstrong and Zhou (2011) show there is a linear relationship between a supplier's price and its commission. Specifically, the lowest retail price a supplier offers is $p_{\min} = (1 - \sigma)v$, and if a supplier chooses retail price p it will offer a salesman the commission payment

$$b(p) = \frac{1 - \sigma}{\sigma}(p - p_{\min}) .$$

As in Varian's original model, one can show that the surpluses of savvy and non-savvy consumers increase with σ , while the total profits of suppliers and salesmen combined decreases with σ .

Using this model one can consider the impact of a policy which restricts the use of

commission payments. Suppose that salesmen remain necessary for consumers to buy the product, but commission payments are banned and a salesman is instead paid directly for a consultation by consumers. Competition between salesmen implies that their consultation charge is zero. Suppose that when a salesman receives no commissions, he steers the naive consumers to the cheaper product. (This might be because, all else equal, he has a small intrinsic preference for assigning the appropriate product to consumers.) In this case, *all* consumers buy the cheaper product and in Bertrand fashion the sellers are forced to set their retail prices equal to cost. It follows that both groups of consumers are better off in the no-commission regime (although salesmen and suppliers are worse off).¹⁵

2.2 The equilibrium number of savvy consumers

When consumers choose to be savvy

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.¹⁶ To discuss the equilibrium extent of savviness, continue with the model of price dispersion from the previous section, 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 we 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 , and where the incentive to become informed, $V_S(\cdot) - V_N(\cdot)$, is "hump-shaped" such that $V_S(0) - V_N(0) = V_S(1) - V_N(1) = 0$.

Suppose that consumers can switch from being ignorant to informed by incurring an information acquisition cost, κ . In general, consumers may differ in their cost of acquiring information, and let $\kappa(\sigma)$ be the corresponding cost of the marginal consumer when σ consumers choose to become informed. A consumer with information acquisition cost κ will choose to become informed if and only if $\kappa \leq V_S(\sigma) - V_N(\sigma)$, and consumers will choose to become informed until the marginal consumer is indifferent. Thus, the fraction

¹⁵Inderst and Ottaviani (2012) present an alternative model of mis-selling, where the salesman advises consumers about the suitability of a product rather than its price. There, no consumers are informed, and must rely on the salesman to advise them about which product to buy. The salesman has only a noisy signal about the suitability of a product, and he has an intrinsic preference to recommend the suitable product to a consumer. However, this preference can be overturned if a seller sets a high enough commission.

¹⁶The issue of how many agents rationally decide to remain uninformed in a market equilibrium was highlighted early on by Grossman and Stiglitz (1980) and Burdett and Judd (1983).

σ of consumers who become informed in an interior equilibrium with $0 < \sigma < 1$ satisfies

$$V_S(\sigma) - V_N(\sigma) = \kappa(\sigma) . \quad (2)$$

Figure 3 illustrates the equilibria, where the hump-shaped curve, $V_S(\sigma) - V_N(\sigma)$, captures the benefit of being informed, while the upward-sloping line $\kappa(\sigma)$ represents the cost of becoming informed. The figure shows there are two interior equilibria satisfying (8), a low- σ and a high- σ equilibrium. However, only the high- 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 any interior equilibrium, because of the search externality too few consumers choose to be informed—and too many prefer to free-ride on other consumers’ search efforts—and aggregate consumer surplus would be boosted if σ were locally increased.¹⁷ (By contrast, if the market instead had a ripoff externality, in the sense that aggregate consumer surplus was a decreasing function of σ , there would be excessive numbers of consumers choosing to be savvy.)




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

If $\kappa(0) > 0$, as depicted on the figure, there is a second stable equilibrium where $\sigma = 0$. When no one becomes informed, all consumers obtain the same (bad) deal in

¹⁷Formally, aggregate consumer surplus when σ consumers incur the cost of being informed is

$$\sigma V_S(\sigma) + (1 - \sigma)V_N(\sigma) - \int_0^\sigma \kappa(\tilde{\sigma})d\tilde{\sigma}$$

which is strictly increasing in σ at any point satisfying (2).

the market, and there is no point investing in acquiring information to find a better deal. This equilibrium is akin to Diamond (1971)’s paradox. However, if a fraction of consumers actively enjoy shopping, so that $\kappa(\sigma) = 0$ for sufficiently all small σ , the unique equilibrium may be the high- σ equilibrium.

One can imagine consumer policies which affect either the cost curve or the benefit curve. Assuming that it is the high- σ equilibrium which is relevant, a policy which reduces information acquisition costs—in the sense of shifting the curve $\kappa(\sigma)$ downwards—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, in the model of price dispersion in section 2.1, we saw on Figure 1 that increasing the number of suppliers 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 suppliers will in equilibrium benefit all consumers once the impact on σ is taken into account.¹⁸

On the other hand, a policy which shifts the benefit curve downwards will reduce the fraction of consumers who choose to become informed.¹⁹ Consider the model of price dispersion discussed in section 2.1, specialized to the case with two sellers, consumer valuation $v = 1$ and costless production as depicted on Figure 1. 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 outlay (including information costs where relevant) of about 0.1. In this example most consumers obtain what seems like a good deal, obtaining the item in return for an 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. Suppose in response that a new policy constrains firms to set prices no higher than $\frac{1}{2}$, so that the maximum permitted price is halved. For a given σ , the expected prices paid by the informed and uninformed consumers then halve, and hence the incentive to become informed also

¹⁸To take an extreme example, if all consumers have information acquisition cost $\kappa = \frac{1}{10}$, then by examining Figure 1 we see that the only equilibrium with duopoly involves no consumers becoming informed, in which case all consumers are charged the monopoly price $p = 1$. However, with four suppliers, the maximum gap between V_S and V_N is greater than κ , and a stable equilibrium with $\sigma \approx 1$ emerges where all consumers have total outlay of about $\kappa = \frac{1}{10}$. A contrasting effect is discussed in Spiegel (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.

¹⁹See Fershtman and Fishman (1994) and Armstrong, Vickers, and Zhou (2009) for analysis of this issue.

halves. The result is that the fraction of informed consumers falls to $\sigma \approx 0.74$, so that the number of uninformed consumers rises about 5-fold as a result of the policy. Each consumer now has expected outlay of about 0.17, which is 70% higher than in the absence of the price cap. Industry profit more than *doubles* as a result of the imposition of the price cap. Thus, the perverse effect of this well-intentioned consumer policy can be substantial.²⁰

When firms confuse consumers

The previous section discussed how consumers can take the initiative to become savvy. Clearly, firms also play role in supplying information to consumers, and there is a vast economic literature about how firms advertise their products and prices. Less familiar is the possibility that the firms attempt to “confuse” consumers, with the result that the fraction of savvy types falls. For example, firms might present their prices in an opaque way or in a different format to their rivals, and this makes it hard for some consumers to compare deals.²¹

To illustrate this possibility, consider the following extension to the price dispersion model of section 2.1.²² There are two firms, and a firm can present its price in one of two formats. If firms choose the same format, consumers find it easy to compare prices and all choose to buy from the firm with the lower price. However, if firms 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 even across formats).

In this context, firms choose both prices and formats according to a mixed strategy. Since the format itself does not matter, only whether the formats are the same or not, a firm chooses the same CDF for its price, say $F(p)$, regardless of its chosen format, and is equally likely to choose either format. If one firm chooses a particular format and price p ,

²⁰Knittel 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 price rises.

²¹For 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.

²²This discussion is based on Piccione and Spiegler (2012) and Chioveanu and Zhou (2013).

then as in expression (1) its expected profit is

$$\left(\underbrace{\frac{1}{2}[1 - F(p)]}_{\text{same format}} + \underbrace{\frac{1}{2}[\sigma(1 - F(p)) + \frac{1}{2}(1 - \sigma)]}_{\text{different format}} \right) (p - c) \equiv \frac{1}{4}(1 - \sigma)(v - c) .$$

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

It is not an equilibrium for firms to choose their format deterministically. Clearly, if both firms were known to choose the same format, price would be driven down to cost and profit to zero. In that case, a firm could switch format to make money from the newly confused consumers. If both firms were known to choose distinct formats, prices would be chosen according to a mixed strategy as in (1). However, in that case, a firm 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.

This model predicts that firms engage in “tariff differentiation” to obtain positive profits, just as firms 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 model, move the market to Bertrand price competition, and all consumers would benefit.²³

2.3 Coasian pricing

Consider next a very different kind of model, the durable goods problem of Coase (1972). There, a firm sells its product over time to forward-looking consumers with heterogeneous tastes for a single unit of 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.

²³ Additional features play a role when the two formats are “simple” and “opaque”, and when both firms choose an opaque format even more consumers are confused relative to when firms choose distinct formats. In such a setting, when a firm 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.

(However, in contrast to the previous model, here price dispersion does not arise through the use of mixed strategies.)

We extend this classical model to allow a fraction $1 - \sigma$ of consumers to be naive, in the sense that they do not understand the firm’s incentive to reduce its price over time.²⁴ As such, these naive consumers buy myopically, as soon as the price falls below their valuation for the item. It follows that $V_N \leq V_S$ since naive consumers buy too soon relative to the optimal purchasing strategy followed by the savvy, and since the firm obtains greater profit when a consumer buys more quickly we have $\Pi_N \geq \Pi_S$. From this perspective, naive consumers here 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, with the result that naive consumers are protected by the presence of the savvy, while profits are harmed. However, in this market savvy consumers might exert their search externality by following an *inefficient* strategy, which is to delay their purchase, and this makes the welfare impact of savvy consumers less clear-cut than it was in section 2.1.

To illustrate these points, consider a simple example. A firm with costless production sells its product over infinite discrete time. All consumers are present from the start and wish to have one unit of the product. There is a binary distribution for consumer valuations: with probability α a consumer has high valuation v_H and with probability $1 - \alpha$ she has lower valuation v_L . Suppose that

$$\alpha v_H \geq v_L , \tag{3}$$

so in a one-period setting the firm prefers to sell only to the high-value consumers than to all consumers. The firm and consumers have discount factor $\delta \leq 1$. A fraction σ of consumers are savvy and foresee the firm’s incentive to reduce its price over time, while the remaining consumers are naive and mistakenly believe the firm’s price will not change and so decide whether to purchase in the initial period myopically.

As soon as all high-value consumers have purchased, the firm will set price $p = v_L$ and sell to the low-value consumers. The price p which just induces the savvy high-value

²⁴For example, when *Apple’s* iPhone was launched in 2007, many early buyers complained when the price fell by \$200 two months after the launch. These consumers might be classified as “naive”, although in this instance they were so vociferous that the company offered them a \$100 voucher as compensation. See the *New York Times* article (September 7, 2007) titled “iPhone owners crying foul over price cut”.

consumers to buy now, anticipating that the price will fall to v_L next period, satisfies $v_H - p = \delta(v_H - v_L)$, so that

$$p = v_M \equiv \delta v_L + (1 - \delta)v_H$$

is intermediate between the high and low valuations.

One can show that the three strategies the firm might follow are:

Strategy 1: Set high initial price $p_1 = v_H$, then medium price $p_2 = v_M$, then low price $p_3 = v_L$.

Strategy 1 involves setting a high price to high-value naive consumers which is not attractive to savvy high-value consumers who anticipate a lower price later. The firm then sets a medium price which is attractive to high-value savvy consumers, and finally sets a low price to mop up all low-value demand. In particular, even though valuations are binary, the firm offers three distinct prices. This strategy generates total discounted profit of $(1 - \sigma)\alpha v_H + \delta\sigma\alpha v_M + \delta^2(1 - \alpha)v_L$, or

$$(1 - \sigma(1 - \delta + \delta^2))\alpha v_H + \delta^2(1 - \alpha + \sigma\alpha)v_L. \quad (4)$$

These profits decrease with σ . When this strategy is used, the naives observe in period 2 that the firm does reduce its price over time, and so might be converted to savvy types. However, by that point the high-value consumers have purchased, and the remaining low-value consumers do not change their behaviour if they do become savvy.²⁵

Strategy 2: Set high initial price $p_1 = v_H$ then low price $p_2 = v_L$.

This strategy yields profits of

$$(1 - \sigma)\alpha v_H + \delta(\sigma + (1 - \sigma)(1 - \alpha))v_L \quad (5)$$

²⁵One advantage of this model is that the naive consumers need be naive only in the initial period, and it makes no difference to the analysis if their “eyes are opened” after the initial period and they are then converted into savvy types. Besanko and Winston (1990) analyze a related model in which consumer valuations are continuously and uniformly distributed and there is a finite time horizon. They compare the most profitable price path when all consumers are forward looking to that when all consumers are non-strategic and buy myopically. They show that the firm chooses a higher initial price with myopic consumers, but the comparison between the final prices is ambiguous. However, if one solves their model with an infinite horizon it seems that the price path for the strategic consumers is uniformly below that for the naive. For instance, if valuations are uniformly distributed on $[0, 1]$, production is costless and the discount factor is $\sigma < 1$, the equilibrium price with naive consumers in period $t = 1, 2, \dots$ is $p_t = (1 + \sqrt{1 - \delta})^{-t}$, while the price with strategic consumers is $p_t = \sqrt{1 - \delta} \times (1 + \sqrt{1 - \delta})^{-t}$. However, it is perhaps implausible that naive consumers continue to be surprised by price reductions after they have seen the firm reduce its price already.

which also decreases with σ . Strategy 1 yields greater profit than strategy 2 if and only if

$$\sigma\alpha(v_H - v_L) \geq (1 - \alpha)v_L . \quad (6)$$

Given assumption (3), this condition is satisfied if and only if σ is large enough.

Strategy 3: Set medium initial price $p_1 = v_M$ then low price $p_2 = v_L$.

When the firm chooses to start with the medium price v_M it will sell to all high-value consumers immediately. The discounted profits with this strategy are $\alpha v_M + \delta(1 - \alpha)v_L$, or

$$(1 - \delta)\alpha v_H + \delta v_L , \quad (7)$$

which do not depend on σ . Condition (3) implies that this profit is higher than v_L , which is the profit if the firm initially charged the low price v_L . Thus there is no need to consider a fourth strategy to sell to all consumers immediately.

A low-value consumer obtains zero surplus in any event. A high-value consumer is worst off when strategy 1 is used, and best off with strategy 3. (High-value naive consumers are indifferent between strategies 1 and 2, while high-value savvy types are indifferent between strategies 2 and 3.) Except when strategy 3 is followed, naive consumers obtain lower surplus than savvy types, since they are too inclined to buy in the first period compared with the optimal purchasing strategy followed by savvy consumers. All else equal, total welfare is also lowest with strategy 1 and highest with strategy 3.²⁶

The firm makes lower profits when the fraction of savvy types is higher. (Its profit is the maximum of the three functions (4), (5) and (7), all of which weakly decrease with σ .) When $\sigma \approx 0$, so almost all consumers are naive, strategy 2 is the most profitable, while when $\sigma \approx 1$ strategy 3 is the most profitable. It follows that strategy 1 can be optimal only with a mixed population of naive and savvy consumers. Clearly, consumers are better off when almost all consumers are savvy compared when almost all are naive.

As we move from $\sigma = 0$ to $\sigma = 1$, it may be that strategy 1 is never followed.²⁷ In this case, strategy 3 is used if and only if σ is large enough, and each consumer's surplus is an

²⁶Provided that the seller's strategy does not change, total welfare weakly decreases with σ . In each of the three strategies, a low-value consumer buys at the same time regardless of whether they are naive or savvy. However, a high-value consumer buys earlier if she is naive than if she is savvy, and this is good for overall welfare. (With strategy 3, all high value consumers buy in the first period and welfare does not depend on σ .)

²⁷The condition for this is $\delta(v_H - v_L)(\alpha v_H - v_L) < (1 - \delta)(1 - \alpha)v_L v_H$, which does not depend on σ . This condition is satisfied if $\alpha v_H \approx v_L$ but violated when δ is close to 1.

increasing function of σ . However, in other cases we move from strategy 2 to strategy 3 via strategy 1.²⁸ Here, savvy consumers are worst off when σ lies in an intermediate range, and their net surplus is “ U -shaped” in σ . Regardless of whether strategy 1 is sometimes used, though, naive consumers who are indifferent between strategies 1 and 2 are always weakly better off as σ increases. As such, this market exhibits search externalities in the classification used in section 1.

In this framework, total welfare is U -shaped in σ . Welfare is the same when $\sigma = 0$ as when $\sigma = 1$, since in either case all high-value consumers buy in period 1 and all low-value consumers buy in period 2. However, for intermediate values of σ strictly fewer high-value consumers buy in the first period when strategy 1 or 2 is followed.

2.4 Quality dispersion

Consider next the supply of a more complicated product with endogenous quality. We can interpret quality quite broadly to include add-on charges and other “small-print” terms. For instance, a seller of insurance may advertise a headline premium, while details about excesses and exclusions are more hidden or hard for some consumers to interpret. Or a snack could be made expensively using good ingredients or made cheaply by using lots of salt, but only a fraction of consumers know how to interpret the list of ingredients.

Specifically, suppose that $n \geq 2$ symmetric firms serve a market. Each firm chooses the price, p , and the quality, q , of its product. All consumers observe the prices from all firms. However, only a fraction σ of savvy consumers also observe all qualities, while the remaining $1 - \sigma$ see no firm’s quality. The less informed consumers are Bayesian, and calculate a firm’s equilibrium incentives to choose quality. All consumers have the same preferences, and their surplus from a product with price p and quality q is $q - p$. We assume consumers are risk-neutral (and in particular, they care about the expected quality of the product if they do not observe quality directly), and their outside option is zero. If a firm chooses quality q , its unit cost of supply is $c(q)$, which is a convex function with $c(0) = c'(0) = 0$. Each firm chooses its price-quality pair (p, q) simultaneously, and simultaneously with its rivals.

²⁸For example, with parameter values $\alpha = \frac{1}{2}$, $v_L = 1$, $v_H = 4$ and $\delta = \frac{2}{3}$ (so that $v_M = 2$), one can check that when $\sigma < \frac{1}{3}$ the firm follows strategy 2, and sets initial price $p_1 = 4$ followed by $p_2 = 1$. For intermediate $\frac{1}{3} < \sigma < \frac{2}{3}$, the firm follows strategy 1, and sets initial price $p_1 = 4$, then medium price $p_2 = 2$, then low price $p_3 = 1$. For $\sigma > \frac{2}{3}$, the firm starts with the medium price $p_1 = 2$ and then drops its price to $p_2 = 1$.

In the extreme cases when $\sigma = 0$ or $\sigma = 1$ the outcome involves pure strategies and zero profits. If $\sigma = 0$, no consumer observes quality and so there is no incentive for a firm supply positive quality, although competition forces firms to set price equal to marginal cost, so that $p = q = 0$. If $\sigma = 1$, each firm maximizes consumer surplus $q - p$ subject to its break-even constraint $p \geq c(q)$, so the efficient quality which maximizes $q - c(q)$ is chosen and price again just covers cost.

However, when $0 < \sigma < 1$ there is no symmetric pure strategy equilibrium. To see this, suppose that all firms choose (p, q) and share the market equally. First, note that $p = q = 0$ is not an equilibrium, since a firm can deviate and offer a higher-quality product at a positive price, sell to savvy consumers, and make a positive profit. Therefore, assume that $p, q > 0$. Then for a firm to have no incentive to cheat (i.e., offer $\tilde{q} = 0$) and serve only its share of the non-savvy we require

$$\frac{1}{n}(p - c(q)) \geq \frac{1}{n}(1 - \sigma)p \Leftrightarrow \sigma p \geq c(q) .$$

In particular, there is a strictly positive mark-up $p - c(q)$ in this candidate equilibrium. However, another possible deviation involves a firm slightly increasing its quality, keeping its price unchanged, which attracts all the savvy consumers. For it to have no incentive to do this we require that

$$\frac{1}{n}(p - c(q)) \geq \left(\sigma + \frac{1}{n}(1 - \sigma)\right)(p - c(q)) \Leftrightarrow n \leq 1$$

which is a contradiction.²⁹

We now derive a symmetric mixed strategy equilibrium.³⁰ In this equilibrium all firms offer the same deterministic price p^* and choose their quality according to a mixed strategy which has an “atom” at $q = 0$ and is continuously distributed for $q \geq p^*$. (If a firm chooses $q < p^*$, no savvy consumer will ever buy and so the firm should “cheat” to the maximum extent and set $q = 0$.) Thus, this equilibrium exhibits quality but not price dispersion. If a firm chooses an unexpected price $p \neq p^*$, uninformed consumers do not buy from it.³¹

²⁹This discussion is adapted from Proposition 2 in Cooper and Ross (1984).

³⁰Details for the following analysis are available from the author on request. Dubovik and Janssen (2012) examine a similar model and issues. However, they assume there are also some totally uninformed consumers who see neither prices nor qualities and buy randomly. When there are enough such consumers, they show there is a mixed strategy equilibrium in which firms choose price according to a mixed strategy, and conditional on its price a firm chooses quality deterministically, so that price is a perfect indicator of quality.

³¹For instance, if a firm chooses a lower price $p < p^*$, uninformed consumers think its quality is zero, while if a firm chooses a higher price $p > p^*$ uninformed consumers believe its average quality is no better than a firm choosing $p = p^*$. In either case, an uninformed consumer does not buy if $p \neq p^*$.

Write the CDF for a firm's choice of quality as $G(q)$, which has support $\{0\} \cup [p^*, q_{\max}]$ where q_{\max} is the highest quality chosen in this equilibrium. Since a firm must be indifferent between choosing any q in this support, for $q \geq p^*$ the CDF G satisfies

$$(\sigma(G(q))^{n-1} + \frac{1}{n}(1 - \sigma))(p^* - c(q)) \equiv \frac{1}{n}(1 - \sigma)p^* ,$$

which is the counterpart to expression (1) above. (The left-hand side shows that the firm attracts its share of the uninformed, and sells to all savvy consumers if its quality is above that of all its rivals. The right-hand side is its profit if it cheats and sets $q = 0$.)

To make further progress, specialise the model to duopoly with a quadratic cost function, so that $n = 2$ and $c(q) = \frac{1}{2}q^2$. In this case, using

$$p^* = \frac{\sigma}{1 + \sigma} \tag{8}$$

in the above construction constitutes a valid equilibrium for any $0 < \sigma < 1$.³² In this equilibrium, the probability that a given firm cheats and offers $q = 0$ is

$$\frac{1 - \sigma}{2(2 + \sigma)}$$

(which decreases with σ), while the maximum quality offered is

$$q_{\max} = \frac{2\sigma}{1 + \sigma} .$$

This maximum quality is below the efficient quality level (which is 1 in this example) and allows a firm to break even (i.e., $c(q_{\max}) \leq p^*$). In this equilibrium, industry profit is

$$\Pi(\sigma) = \frac{\sigma(1 - \sigma)}{1 + \sigma} \tag{9}$$

which is zero at each extreme $\sigma = 0$ and $\sigma = 1$. (We have already seen that there is Bertrand price competition in these cases.) Thus, unlike the models of price dispersion discussed earlier, here suppliers make low profits when most consumers are non-savvy, since there is nevertheless competition in terms of price p which acts to dissipate profit. Since savvy consumers buy the product with the higher quality, firms extract less profit from them than from an uninformed buyer, and $\Pi_N > \Pi_S$.

The following figure plots the net surplus of each savvy consumer (as the middle curve), the net surplus of each uninformed consumer (as the lower curve), and total welfare (consumer surplus plus industry profit, plotted as the upper curve in bold). Thus, as in the

³²In fact, there is an interval of prices p^* which constitute this kind of "fixed price" equilibria, and the price in (8) is the smallest such price. At the extremes where $\sigma \approx 0$ or $\sigma \approx 1$, the range is very narrow.

model of price dispersion in section 2.1, each savvy consumer confers a positive search externality on other consumers. When most consumers are savvy and know the qualities offered by suppliers, suppliers compete to offer an efficient combination of price and quality, which the uninformed can usually enjoy too. When most consumers cannot discern quality, a supplier has little incentive to offer high quality, and even a savvy consumer is unlikely to secure a good product in such a market.

The model predicts that competitive firms set a rigid price, but differ in their quality which is only observed by the savvy buyers. Rational but uninformed buyers are put off by a seller which offers a lower price, and presume that such a seller will be cheating on quality. A market which might fit the model is insurance, where a particular kind of insurance might be offered by rival sellers at a similar price, but different sellers might have more “exclusions” than others which only the savvy can notice and avoid.³³




Figure 3: surplus and welfare with quality dispersion

This model, and others like it, is considerably easier to analyze if the non-savvy consumers are strategically naive, and do not make a connection between price and quality. For instance, suppose that non-savvy consumers anticipate quality q^e regardless of the price asked, and so buy from the firm with the lowest price (provided the price is below q^e). In a competitive market with many (at least four) firms, an asymmetric equilibrium in

³³Recall section 2.2, where the imposition of a price cap acted to raise average prices. One could try to perform a similar exercise in the context of quality dispersion. That is, a policy which imposes a minimum quality standard for products supplied in the market might reduce the incentive for consumers to choose to become informed about quality, with the net result that average quality falls.

pure strategies exists of the following form: some firms serve savvy consumers by offering the efficient level of quality at a price which just covers cost, while other firms serve naive consumers by offering quality $q = 0$ at the price which just covers the associated cost (i.e., $p = 0$ since $c(0) = 0$). In such a market, the proportion of savvy types, σ , has no impact on the deals offered to either type of consumer, and savvy types do not protect or harm the interests of the rest.³⁴ We discuss a closely related model of add-on pricing in section 3.2 in more detail.

3 Hold-up

A market exhibits “hold up” when consumers are, to some extent, committed to purchasing the product before they know the full terms of trade. For instance, if a consumer must make a costly journey to a seller to discover its price, she might decide to buy even if the price she finds there was somewhat higher than she anticipated. Likewise, some consumers may have to decide whether to purchase a product without being able to discern its quality, or without knowing the prices of “add-on” products which later become available. We discuss how these markets perform in situations with an indivisible good (section 3.1) and in the more complex case with add-on pricing (section 3.2).

3.1 An indivisible product

The model of price dispersion in section 2.1 involved a market where most consumers paid high prices when there were only a few informed consumers present. In hold-up situations, Diamond (1971) shows how the market can break down altogether. In this section, we see how the presence of savvy consumers can overcome or amplify this danger, depending on the precise form of savviness.³⁵

Suppose that a single supplier sells a product with unit cost c to a population of consumers. Consumers differ in their reservation value for the item, v , where the fraction of consumers with $v \geq p$ is denoted $q(p)$. Suppose for now that all consumers know their value v in advance. Crucially, all consumers must incur a travelling cost $t > 0$ to reach

³⁴Armstrong and Chen (2009) analyze a related model where mixed strategies are used, and find a symmetric equilibrium in which firms offer random prices and obtain positive profits, and where price is a perfect indicator of quality, i.e., if a firm chooses a high price it will also choose high quality. However, the naivete of the uninformed consumers prevents them acting on this indicator. In this equilibrium, the fraction of savvy types does affect the surplus obtained by the savvy and naive consumers.

³⁵The discussion in this section is related to Stiglitz (1979) and Anderson and Renault (2006).

the seller. Thus, the type- v consumer will choose to travel to the seller if $v \geq t + \tilde{p}$, if she observes (or expects to pay) price \tilde{p} . Suppose a fraction σ of consumers (independent of v) are savvy and know the seller's true price in advance (but still incur the cost t if they choose to buy), while the remaining $1 - \sigma$ have to travel to the seller to discover its price. These $1 - \sigma$ ignorant consumers are rational, though, and anticipate the seller's incentive to set its price.

The equilibrium price is derived as followed. Suppose that p^* is the price that an uninformed consumer expects to pay. If the seller actually chooses the price p , where p is not too much bigger than p^* in the sense that $p \leq p^* + t$, its demand is

$$\sigma q(p + t) + (1 - \sigma)q(p^* + t) .$$

The informed travel to the seller (and buy) if $v \geq p + t$, while the uninformed travel to the firm if $v \geq p^* + t$ and once at the seller they buy provided that $v \geq p$. Thus, similarly to the model in section 2.1, the demand from the uninformed is inelastic, at least with respect to local changes around p^* . Since the firm is free to choose its price p given the price anticipated by the uninformed consumers, an equilibrium price p^* is such that choosing price $p = p^*$ must

$$\text{maximize}_{p \leq p^* + t}: (\sigma q(p + t) + (1 - \sigma)q(p^* + t))(p - c) .$$

This problem has first-order condition

$$q(p^* + t) + \sigma(p^* - c)q'(p^* + t) = 0 . \tag{10}$$

If the demand function $q(\cdot)$ is log-concave, this first-order condition has a single solution which determines the equilibrium price p^* .³⁶

If $\sigma = 1$, then the equilibrium price is the price that maximizes profit $(p - c)q(p + t)$. If $\sigma = 0$, though, so no consumers know the price in advance, no consumer chooses to travel to the seller, and the market breaks down altogether. The seller knows that every consumer is willing to pay t more than their anticipated price p^* for the item, and so it has an incentive to set its price at least equal to $p^* + t$, and there is no equilibrium price which induces any consumer to incur the travel cost t . When some consumers are informed in

³⁶Note that the limit of this price as $t \rightarrow 0$ is not the price which the firm would charge if there was no travel cost and all consumers were willing to travel to the firm to find out its price. (This price would be the monopoly price which maximizes $(p - c)q(p)$.) Thus, there is a discontinuity in the outcome between $t = 0$ and arbitrarily small $t > 0$.

advance, though, the market opens, which is to the advantage of both consumers and the seller.

When $q(\cdot)$ is log-concave, formula (10) implies that the equilibrium price is a decreasing function of σ . Therefore, all consumers—savvy and uninformed—benefit when σ rises. Since the equilibrium price p^* is above the monopoly price that maximizes $(p - c)q(p + t)$, and this profit $(p - c)q(p + t)$ is single-peaked in p , it follows that the firm too is better off when σ is larger. Thus, this market provides an example where the search externality is present, and where boosting the fraction of savvy types benefits all parties.³⁷ (By contrast, in section 2.1 industry profits were decreasing in the fraction of informed consumers.)

Contrasting effects are seen if some consumers know their value v in advance rather than the price. Suppose that no consumer knows the price in advance, so the danger of hold-up is present. However, a fraction σ of consumers are savvy in the sense that they know their value v in advance, while the remaining $1 - \sigma$ consumers only discover their valuation once they travel to the seller and inspect the product. (These uninformed consumers view the distribution of uncertain values to be governed by the function q .)

An uninformed consumer who expects to pay price p^* will travel to the seller if expected surplus is greater than their travel cost, i.e., if $t \leq s(p^*)$, where

$$s(p) \equiv \int_p^\infty q(\tilde{p})d\tilde{p} , \quad (11)$$

is net consumer surplus (the “area under the demand curve”) with price p . Informed consumers, by contrast, will travel to the seller if $t \leq v - p^*$. Provided that $s(p^*) \geq t$, the seller’s demand when it chooses price $p \leq p^* + t$ is

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

since all uninformed consumers travel to the seller, and they will buy if they discover that $v \geq p$. Thus, now the *informed* consumers have locally inelastic demand. An equilibrium price p^* is such that choosing $p = p^*$ must

$$\text{maximize}_{p \leq p^* + t}: (\sigma q(p^* + t) + (1 - \sigma)q(p)) (p - c) ,$$

which has first-order condition

$$\sigma q(p^* + t) + (1 - \sigma)q(p^*) + (1 - \sigma)(p^* - c)q'(p^*) = 0 . \quad (12)$$

³⁷As in section 2.2, therefore, profits can be increased if a price cap is imposed, albeit for a very different reason. In the current context, a price cap operates to commit the firm not to set high prices.

Following similar arguments to those used when some consumers knew the price in advance, when q is log-concave, the first-order condition for this problem uniquely determines the candidate equilibrium price, but now this price p^* *increases* with σ . When $\sigma = 0$, the candidate price is the monopoly price p^M which maximizes $(p - c)q(p)$, while when $\sigma = 1$ the candidate price is such that $q(p^* + t) = 0$. If the requirement $s(p^*) \geq t$ fails, the uninformed have no incentive to participate, and when this happens the market shuts down. When the travel cost t is so large that $t > s(p^M)$, the uninformed will not travel to the firm even in the most favorable case when $\sigma = 0$. However, if $t < s(p^M)$, then the market opens if σ is sufficiently small (and fails when σ is close enough to 1).

In the range of σ where the market opens, the equilibrium price increases with σ , and all consumers as well as the firm are worse off with larger σ . When the market is open, savvy consumers obtain a higher surplus than the uninformed, since a consumer would prefer to know her valuation before deciding to travel to the firm. (Expected surplus of a savvy consumer is $V_S = s(p^* + t)$, while expected surplus of an uninformed consumer is $V_N = s(p^*) - t$, which is smaller.) Since demand from the uninformed, $q(p^*)$, is higher than from the informed, $q(p^* + t)$, the firm obtains more profit from the non-savvy, and $\Pi_N > \Pi_S$.

To illustrate, in the linear demand example where $q(p) = 1 - p$ and $c = 0$, the price which solves (12) is

$$p^* = \frac{1 - t\sigma}{2 - \sigma} . \quad (13)$$

At this price, the requirement that $s(p^*) \geq t$ is never satisfied when $t \geq s(p^M) = \frac{1}{8}$, in which case the market shuts down. If $t < \frac{1}{8}$, however, the market is open if the fraction of informed consumers σ is sufficiently small.³⁸ In this example, aggregate consumer surplus V and profits Π both decrease with σ , and hence welfare does too.

Using the terminology of section 1, in formal terms this market exhibits a ripoff externality.³⁹ Uninformed consumers may be willing to invest in travelling to the seller to discover their valuation, for the chance they like the product, and this helps the market remain open.⁴⁰ However, despite the fact that $\Pi_N(\sigma) > \Pi_S(\sigma)$, it is perhaps misleading to

³⁸The precise condition is that $\sigma \leq \frac{1-2\sqrt{2t}}{1-\sqrt{2t-t}}$.

³⁹Using similar analysis to that in section 2.2 one could investigate the equilibrium proportion of savvy consumers, when a consumer can choose to become informed of her valuation *ex ante* by incurring a cost κ . Because a ripoff externality operates in this market, we expect that *too many* consumers choose to become informed in equilibrium.

⁴⁰See Anderson and Renault (2000) for related analysis in a duopoly model. In that model, the price-

say that the uninformed consumers are being “ripped off”. Unlike the market described in the next section, the seller here is not engaging in any tactic which aims to exploit this group of consumers.

3.2 Add-on pricing

Hold-up can also occur when a seller supplies an “add-on” product once a consumer has purchased the 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; the ability to obtain a casual overdraft from your bank without prior agreement, or the ability to have your luggage stored in the aircraft’s hold in the event it is deemed slightly too large for the cabin.

In some of these examples, it may be that the firm does not choose its add-on price until the customer has purchased the core product, in which case the firm is tempted to set monopoly prices for these services. In other cases, though, the firm chooses both prices at the same time, and the issue is not one of lack of commitment. Rather, the problem is that 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. In this section we explore these cases where firms choose both prices simultaneously, but some consumers cannot, or do not, take adequate notice of the add-on price. Three variants are discussed: the first where non-savvy consumers are rational, but cannot see or interpret the add-on price; a second where naive consumers do not foresee their demand for the add-on service, and a third where naive consumers can be tricked into paying for add-ons they don’t want. Perhaps surprisingly, these apparently small changes in model assumptions generate all three of the market scenarios listed in section 1.

Rational but uninformed consumers: Here we assume that some consumers have prohibitive costs for reading and/or understanding the “small-print” in the contract to discover terms for add-ons, although they do care about these terms. (More generally, the following discussion is isomorphic to a model in which firms choose the quality of their product, and only a fraction of consumers are able to discern quality directly.)

raising impact of savvy consumers acts to boost industry profit, in contrast to the monopoly case where the equilibrium price is too high from the seller’s perspective.

A single seller supplies a core product, for which consumers have heterogeneous value X . The price of this product is denoted P and its unit cost is C . The fraction of consumers with $X \geq P$ is denoted $Q(P)$. Once a consumer has purchased this core product, an add-on product becomes available. The price of this product is p and its unit cost is c . All consumers have the same add-on demand, and with price p a consumer will consume $q(p)$ units of the add-on service.⁴¹ 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. The expected net surplus from the option of being able to buy the add-on at price p is $s(p)$ in (11). Thus, if a consumer with valuation X anticipates (or observes) the add-on price \tilde{p} , she will buy the core product if

$$X + s(\tilde{p}) \geq P . \quad (14)$$

If she does buy the core product, she will go on to generate add-on profit $\pi(p)$, where p is the firm's *true* add-on price.

Suppose a fraction σ of consumers observe the firm's true add-on price, while the remaining $1 - \sigma$ do not. If the uninformed expect to pay add-on price p^* , the firm's expected profit from the two groups of consumers is

$$(\sigma Q(P - s(p)) + (1 - \sigma)Q(P - s(p^*))) \times (P - C + \pi(p)) . \quad (15)$$

Let (P^*, p^*) denote the equilibrium pair of prices. If we assume passive conjectures, uninformed consumers anticipate the add-on price p^* even if they observe an unexpected core price $P \neq P^*$.⁴² For these prices to constitute an equilibrium, choosing $(P, p) = (P^*, p^*)$ maximizes (over any pair of prices P and p) the expression (15). The two first-order conditions for this problem are

$$Q(P^* - s(p^*)) + Q'(P^* - s(p^*)) \{P^* - C + \pi(p^*)\} = 0 ;$$

$$Q(P^* - s(p^*))\pi'(p^*) + \sigma q(p^*)Q'(P^* - s(p^*)) \{P^* - C + \pi(p^*)\} = 0 .$$

⁴¹This 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, and is independently distributed from X .

⁴²When some consumers see one dimension of a seller's choice but not another, this raises the issue of how such a consumer forms her expectation of the unobserved variable given what she does observe. If she observes the firm's core price but not its add-on price, what does she infer about the likely add-on price from the core price? We take the simplest approach and suppose that the uninformed consumers have "passive conjectures", in that they hold beliefs about the firm's add-on price which must be fulfilled in equilibrium, but these beliefs are not affected if a firm deviates from the anticipated core price.

Eliminating terms in P^* reveals that the add-on price satisfies $\pi'(p^*) = \sigma q(p^*)$, or

$$(1 - \sigma)q(p^*) + (p^* - c)q'(p^*) = 0, \quad (16)$$

which has a similar form to the earlier expression (10). Thus, 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 $\pi(p)$. More generally, when add-on demand $q(\cdot)$ is log-concave, formula (16) implies that the add-on price is a decreasing function of σ . For example, when $q(p) = 1 - p$ and $c = 0$, the add-on price is $p^* = \frac{1-\sigma}{2-\sigma}$.

It is straightforward to show that consumers and the firm are worse-off when the equilibrium add-on price is higher, i.e., when σ is smaller.⁴³ Thus, this market exhibits search externalities of the strong kind where *all* parties are better off when the fraction of informed consumers rises. In this it is like the hold-up model with an indivisible product discussed above.

This discussion implicitly assumed that the seller had to offer the same add-on terms to all its customers, which seems a reasonable assumption in most contexts. (It is hard to imagine a hotel supplying rooms with different minibar prices, for instance.) If feasible, though, the seller has an incentive to set different add-on terms: an efficient price $p = c$ aimed at the savvy, and a monopolistic price $p = p^M$ aimed at the non-savvy. (The non-savvy might accidentally choose the efficient contract, but this probability could be reduced if the seller somehow offered the monopolistic contract, together with a “hard-to-find” efficient contract which only the savvy could locate.) The fact that the seller offers the same deal to all consumers immediately implies that $V_S = V_N$ and $\Pi_S = \Pi_N$.

This framework with rational but uninformed consumers is hard to extend to competitive environments, except in the extreme cases where $\sigma = 0$ or $\sigma = 1$.⁴⁴ The model with quality dispersion in section 2.4 has this flavour, and could presumably be re-interpreted as a model of add-on pricing after suitable adjustments. Much easier to analyze, and arguably

⁴³To see this, write $Y = P - s(p)$ for the “total price” for the core product. Then consumers are better off when Y is smaller. For given equilibrium add-on price p^* , the firm chooses the core price P to maximize $Q(P - s(p^*))(P - C + \pi(p^*))$, i.e., it chooses the total price Y to maximize $Q(Y)(Y - C + w(p^*))$, where $w(p) \equiv s(p) + \pi(p)$ is total add-on surplus with add-on price p . The function $w(p)$ is decreasing for $p \geq c$. Since Y maximizes $Q(Y)(Y - C + w)$, a higher w (i.e., a lower p^*) is like the monopolist having a lower cost, which induces a lower Y and higher optimal profits.

⁴⁴Ellison (2005) presents a Hotelling-style duopoly model of add-on pricing. 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 $\sigma = 1$ and $\sigma = 0$ respectively. In his model, industry profits are higher when no consumer is informed of add-on prices, in contrast to the monopoly case just presented.

more applicable in some contexts, are variants when the non-savvy are strategically naive rather than rational. We discuss two such variants in the remainder of this section.

Naive consumers do not foresee need for add-on product: This version of the add-on pricing problem supposes that non-savvy consumers simply 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 more detail, the savvy consumers continue to purchase the core product according to the optimal rule in (14), while naive consumers purchase myopically, i.e., when $X \geq P$. These naive consumers under-estimate the value of the core product, and buy too rarely. (For example, a naive consumer when choosing a hotel room overlooks the benefits of having the minibar in the room.) Nevertheless, once they have purchased the core product, they go on to generate profits $\pi(p)$ for the supplier. This framework can be analyzed in a monopoly or competitive context. The latter is somewhat more transparent, and we focus on that case.⁴⁵

Suppose that many (more than four) sellers compete for consumers by offering a pair of prices (P, p) , where P is a seller's price for the core product and p is its price for the add-on. Savvy consumers foresee their likely future demand, and so buy from a seller with the lowest "total price" $P - s(p)$, provided their valuation X is above this total price. Naive consumers buy from a seller with the lowest core product price P , provided their X is above this price.

An asymmetric equilibrium with pure strategies takes the following form. Some sellers offer an efficient contract aimed at savvy consumers, which has $(P, p) = (C, c)$. Other

⁴⁵The case of monopoly is analyzed as follows. As in the previous version of the add-on pricing problem, suppose that the firm offers the same contract, denoted (P, p) , to all its customers. Similarly to (15), the seller's profit is then $(\sigma Q(P - s(p)) + (1 - \sigma)Q(P)) \times (P - C + \pi(p))$. Note that the profit from a savvy consumer here is *greater* than that from a naive consumer, i.e., $\Pi_S > \Pi_N$, since naive consumers are less inclined to buy the core product. (This contrasts with our analysis of the Coase model in section 2.3, where naive consumers also purchased myopically but where $\Pi_S < \Pi_N$, since in the Coase model the two products were substitutes, while in this add-on model the core product and the add-on are complements.) As such, the firm has an incentive to "educate" its customers, if feasible, in the sense that it wishes to convince naive consumers that they will, in fact, gain value from the add-on service.

One can check that the first-order condition for the optimal add-on price is

$$\frac{\pi'(p^*)}{q(p^*)} = \frac{\sigma Q'(P^* - s(p^*))}{\sigma Q'(P^* - s(p^*)) + (1 - \sigma)Q'(P^*)}$$

which is a less neat formula than the "rational" version in (16) which entailed $\pi'/q = \sigma$. However, in the special case where core product demand Q is linear, the two first-order conditions coincide, and the equilibrium add-on price is exactly the same with rational and with naive consumers.

sellers offer an inefficient contract aimed at the naive which takes a “bargain-then-ripoff” (or “loss-leader”) form. Specifically, the naive contract has a monopoly add-on price p^M and a subsidized core product price which just enables a firm to break even, so that $P - C + \pi(p^M) = 0$. This contract is not attractive to savvy forward-looking consumers, who prefer the cost-reflective tariff, but is attractive to the naive who do not foresee their future demand for the add-on. There is therefore price dispersion in the market, both for the core product and for the add-on. These contracts do not depend on the fraction of savvy types present in the market, and there are no externalities between the two groups of consumers.

Similarly to (11), write

$$S(P) \equiv \int_P^\infty Q(\tilde{P})d\tilde{P}$$

for consumer surplus from the core product alone when its price is P . Then the average surplus of a savvy consumer in this market is $V_S = S(C - s(c))$, while the “true” surplus of a naive consumer with add-on price p is

$$V_N = S(C - \pi(p)) + Q(C - \pi(p))s(p) . \quad (17)$$

(To understand this expression, note that the core product price is $P = C - \pi(p)$ when the add-on price is p , and so a consumer buys if $X \geq C - \pi(p)$, which yields core product surplus $S(C - \pi(p))$. A naive consumer who buys the core product will, in fact, go on to consume the add-on product, and obtain surplus from this product equal to $s(p)$, which explains the second term in the above expression.) The convexity of $S(\cdot)$ together with the fact that $s(c) \geq s(p) + \pi(p)$ for any price p implies that $V_S \geq V_N$.

This model predicts that naive consumers end up paying high add-on prices. As in section 2.2, regulators might consider controlling suppliers’ freedom to exploit consumers in this fashion. However, there are two market failures operating in the *laissez-faire* market—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 may have mixed effects on the naive consumers. (This regulation has no impact on the surplus enjoyed by savvy consumers.) In technical terms, naive surplus in (17) is not necessarily decreasing in p , and regulation to limit “ripoffs” might harm rather

than benefit these consumers.^{46,47}

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. Unless there are strong scale economy effects (so that having large numbers of credulous consumers 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.⁴⁸ 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 agents have self-control problems, and where savvy agents foresee this in advance and naive agents do not.⁴⁹

These various situations all share the same basic structure. The market is competitive, and so supplier profits are driven to zero. All consumers ultimately exhibit the same behaviour and cause the same costs when faced with a given contract, and so the set of contracts which are consistent with zero profits are the same for a naive as for a savvy consumer. From this set, suppliers in equilibrium choose the contract which is most attractive *ex ante* to the target consumer, and so by construction, neither type is tempted

⁴⁶To illustrate, suppose that $Q(P) = 1 - P$, $q(p) = 1 - p$, $C = \frac{3}{4}$ and $c = 0$. The tariff aimed at naives in the unregulated market can be calculated to be $P = p = \frac{1}{2}$, which induces naive surplus $V_N = \frac{3}{16}$. If regulation forced the add-on price down to cost, so that $p = 0$, the core product price rises to $P = \frac{3}{4}$, and naives have reduced surplus $V_N = \frac{5}{32}$.

⁴⁷Other natural ways to model the add-on pricing problem do not have this feature, however. For example, suppose that (very) naive consumers do 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. Once they purchase the printer, though, they realize they do face on-going add-on charges, and their perceived add-on surplus falls from $s(0)$ to $s(p)$ with add-on price p . In this situation, naive consumers buy the core product too often, rather than too rarely, and this makes a reduction in the add-on price unambiguously beneficial to these consumers. In technical terms, the naive surplus in this situation is modified from (17) to $V_N = S(C - \pi(p)) + Q(C - \pi(p))(s(p) - s(0))$, which is always decreasing in p for p between cost c and the monopoly price p^M .

⁴⁸See Sandroni and Squintani (2007) for a model along these lines.

⁴⁹See DellaVigna and Malmendier (2004) and Spiegler (2011, section 2.3) for further discussion.

by the contract aimed at the other group. The outcome is therefore as if a consumer's savviness or naivete was known to suppliers, and there is no interaction between the two groups.

Bill shock: The previous variant could be viewed as a model of “hidden benefits”, in the sense that naive consumers under-estimated how much they will value the service. The final variant we consider, by contrast, is model with “hidden costs”. Here, naive consumers mistakenly buy an add-on service which they do not particularly want or need, a phenomenon sometimes known as “bill shock”.⁵⁰ In this situation, the core price is subsidized with the profits generated by the fraction of naive consumers who end up paying for unwanted add-ons, and this benefits the savvy consumers.

Examples of the kind of add-on “service” we 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.⁵¹ 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.⁵² 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 in a contract with high overage charges.⁵³

⁵⁰Gabaix 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 firm decides whether to advertise or to “shroud” its add-on price. When a firm advertises its price, this acts as an “eye-opener” and consumers realize they will need to pay the charge unless they take evasive action in advance. If firms 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 firms shroud their add-on price, and naive consumers end up paying it.

⁵¹At 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).

⁵²Armstrong 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 having to pay these charges beforehand.

⁵³See Grubb (2014) for analysis of this form of bill shock, and the possible interventions to overcome

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 he 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 know the warranty has no value, or are otherwise immune to the salesman’s patter, and do not buy. In addition, such practices 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 a new supplier), 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.⁵⁴

To model these situations, consider the following stylized framework. Two or more firms 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 extra payment $R > 0$ to their chosen seller once they have purchased the core product. This extra payment might be generated via small-print “traps” or worthless add-ons, which savvy consumers know how to avoid. If sellers cannot distinguish the two kinds of consumers in advance, the equilibrium outcome in this market is for the core product to be subsidized by the anticipated rents from the naive, so that

$$P = C - (1 - \sigma)R . \tag{18}$$

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.

As before, a consumer has idiosyncratic valuation X for the product, $Q(P)$ is the proportion of consumers with $X \geq P$ and $S(P)$ measures consumer surplus from the product (without ripoffs) when price is P . The average surplus of a savvy consumer in

the problem.

⁵⁴Interestingly, the prominent 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).

this market is then $V_S(\sigma) = S(C - (1 - \sigma)R)$, which decreases with σ and so this market exhibits ripoff externalities.⁵⁵ 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. In general 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 . However, 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. However, the impact of such regulation on the two groups of consumers differs: a savvy type benefits from a firm's ability to ripoff the naive and so would like R to be 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 might lobby against this welfare-enhancing regulation.⁵⁶

4 Conclusions

This paper has explored how the balance of “savvy” and “non-savvy” consumers in a market affects firm behaviour and the deals offered to consumers. We discussed two ways in which the two groups might interact: the case of search externalities, where savvy consumers help the non-savvy to obtain a good deal, and the case of ripoff externalities, where non-savvy consumers enable the savvy to obtain a good deal.

We restricted our attention to two broad kinds of market: those which exhibit price or quality dispersion (section 2) and those involving forms of hold-up (section 3). All of the markets examined in section 2 involved a search externality, and the presence of savvy consumers protected the interests of non-savvy consumers. In the model of oligopoly price

⁵⁵It is straightforward to analyze the monopoly version of this market. If the monopolist chooses price P , its profit is $Q(P) \times (P - C + (1 - \sigma)R)$. Thus, the most profitable price is a decreasing function of $(1 - \sigma)R$, so that savvy consumers, as well as the firm itself, are better off when σ is smaller or R is larger.

⁵⁶If R is large enough, the price in (18) 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 the core bank account is typically free.) Firms then have no way to dissipate profits, and profits will be positive even in a competitive market. In these cases, firms as well as savvy consumers have an incentive to lobby against constraints on ripoffs.

dispersion in section 2.1, this was because 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 profits fell with σ while overall welfare increased. Because profits fall with the fraction of savvy types, firms 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 (section 2.2).

The model with quality dispersion in section 2.4 behaved in a broadly similar manner. There, savviness reflected a consumer's ability to discern product quality, and greater numbers of these consumers increased a firm's incentive to offer an appropriate level of quality, which the non-savvy could also enjoy. The fact that the uninformed were nevertheless rational, and understood a seller's incentive to cut quality and supply only the uninformed, implied that the equilibrium involved a "quality assuring" rigid price above cost. In this market, then, there was quality, but not price, dispersion. The savvy obtained higher surplus, and generated lower profits, than the uninformed consumers. A major contrast to the model with price dispersion, though, was that industry profit was non-monotonic in σ , and profits were low even when few savvy types were present. This was because firms have two strategic variables, price and quality, and all consumers observe prices. As such, there was vigorous competition when most consumers were non-savvy.

In the Coasian variant in section 2.3, there was intra-firm price dispersion, and savvy consumers who understood the seller's incentive to reduce price over time intensified the firm's incentives to compete against its earlier self, which benefitted the naive. However, because savvy consumers exerted their positive externality by following a socially *inefficient* strategy—namely, waiting to buy later—their presence could reduce welfare overall.

The models of hold-up in section 3 presented a more mixed picture, and small changes in the model could swing the market from one with search externalities to one with ripoff externalities. This was seen most transparently in section 3.1, where a monopolist supplied an indivisible product and consumers incurred a sunk cost to travel to the seller. If savvy consumers knew the firm's price in advance, while non-savvy consumers had to travel to the seller to discover its price, the former were more price-sensitive. Thus, as in section 2.1, the price decreases with σ , and the search externality benefits the non-savvy and boosts overall welfare. However, with hold-up the supplier cannot refrain from setting a high price to exploit uninformed consumers, and the equilibrium price is too high even from the

seller's perspective. As such, the price reduction caused by savvy consumers increased the seller's profits in this market.

However, if a savvy consumer knew her valuation for the product in advance, while non-savvy consumers had to travel to the seller to see how much they liked the product, the latter were the consumers with more elastic demand. (Neither group of consumers knew the price in advance, and so the savvy consumers had demand which was inelastic with respect to small price rises. The uninformed, however, travelled to the seller in any event, to see if they liked the product, and so their demand could respond to out-of-equilibrium price changes.) In this case, price was an *increasing* function of σ . The presence of non-savvy consumers protected the interests of the savvy, i.e., there was what we termed a ripoff externality. Indeed, too large a fraction of savvy types caused the market to break down altogether.

In section 3.2, three variants of a market with add-on pricing were discussed. Consumers initially decide whether to buy a core product, and if they did so, they were subsequently offered a complementary product. The first variant involved a monopoly seller, and savvy consumers knew the add-on price in advance. Non-savvy consumers did not know this price when deciding to buy the core product, but were rational and understood the firm's incentive to hold them up with a high add-on price. As with the model in section 2.4, the presence of savvy types gave the seller an incentive to set reasonably efficient add-on terms, and the market involved a search externality.

In a second variant, the non-savvy were strategically naive in the sense that they did not foresee their future demand for the add-on product. This myopic perspective led them to purchase the core product too rarely. In a competitive market, there was dispersion in both the core price and the add-on price. Some sellers offered an efficient cost-based tariff aimed at the savvy who understood they would need the add-on service, while other sellers offered a "bargain-then-ripoff" tariff with a subsidized core product price and monopolistic terms for the add-on. This contract was more attractive to naive consumers, since they myopically believed the low price for the core product was good value. In this market, there were no externalities between the two groups of consumers in either direction. Policy which constrained a seller's ability to set high add-on prices had ambiguous effects on the surplus of naive consumers. (It had no impact on savvy surplus.) High add-on prices fund the subsidized price for the core product, and this helps overcome the market failure caused

by these consumers buying the core product too rarely.

The final model involved “bill shock”, where naive consumers could be tricked into making extra payments, while savvy consumers could defend themselves against these tactics. (These tricks might involve penalty charges for errant behaviour, persuasion to buy useless add-ons, or rollover contracts which savvy consumers know to cancel.) Competing sellers set the price for the core product in anticipation that a consumer might be naive and generate extra revenue. As such, the core price was subsidized, with a greater subsidy when the fraction of non-savvy types was larger. Savvy consumers therefore 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 will harm the savvy types who prey on the naive consumers when they fall into small-print traps.

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 Y. CHEN (2009): “Inattentive Consumers and Product Quality,” *Journal of the European Economic Association*, 7(2-3), 411–422.
- 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), 368–395.
- BAYE, M., AND J. MORGAN (2002): “Information Gatekeepers and Price Discrimination on the Internet,” *Economics Letters*, 76(1), 47–51.
- BESANKO, D., AND W. WINSTON (1990): “Optimal Price Skimming by a Monopolist Facing Rational Consumers,” *Management Science*, 36(5), 555–567.
- BROWN, J., AND A. GOOLSBEE (2002): “Does the Internet Makes 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.
- COOPER, R., AND T. ROSS (1984): “Prices, Product Qualities and Asymmetric Information: The Competitive Case,” *Review of Economic Studies*, 51(2), 197–207.
- 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.
- DUBOVIK, A., AND M. JANSSEN (2012): “Oligopolistic Competition in Price and Quality,” *Games and Economic Behavior*, 75(1), 120–138.
- 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. (2014): “Consumer Inattention and Bill-Shock Regulation,” *Review of Economic Studies*, forthcoming.
- INDERST, R., AND M. OTTAVIANI (2012): “Financial Advice,” *Journal of Economic Literature*, 50(2), 494–512.
- 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.
- 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.
- SPIEGLER, R. (2011): *Bounded Rationality and Industrial Organization*. Oxford University Press, Oxford, UK.

STIGLITZ, J. (1979): “Equilibrium in Product Markets with Imperfect Information,”
American Economic Review, 69(2), 339–345.

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