Economic models of consumer protection policies

Armstrong, Mark

Department of Economics, University of Oxford

November 2011

Economic Models of Consumer Protection Policies*

Mark Armstrong
Department of Economics
University of Oxford

November 2011

Abstract

This paper summarizes some of my recent work on consumer protection. I present three theoretical models which illustrate the merits and drawbacks of a number of common consumer protection policies, namely: policies which prevent firms from setting unduly high prices; policies which prevent firms requiring on-the-spot decision making by prospective customers, and policies which prevent suppliers from paying commission payments to sales intermediaries.

1 Introduction

The objective of both consumer and competition policy is to deliver well-functioning markets, something which requires both a strong supply side (competition) and a strong demand side (consumers). For many products, vigorous competition is the single best protection for consumers, and only minimal consumer protection (general contract law, forbidding deceptive marketing, the ability to return faulty goods, and so forth) is needed. As Muris (2002), a former Chairman of the FTC, writes: "[R]obust competition is the best single means for protecting consumer interests." However, in some markets some consumers do not always obtain a good deal, even when substantial competition is present, and in such cases additional policies to aid consumers have a role to play.

*This paper was prepared for the conference on "The pros and cons of consumer protection", organized by the Swedish Competition Authority, held in Stockholm on 11 November 2011. I am very grateful to my discussant, Russell Damtoft, for his insightful comments.

What prevents markets from delivering good outcomes to consumers? Familiar reasons include abuse of dominance and collusion between suppliers, and these fall broadly within the domain of competition policy. However, there are several other reasons why competition need not work well, such as imperfect information about product attributes, imperfect information about market prices, supplier costs of advertising, consumers possessing imperfect information about their own needs, or the use of high-pressure and misleading sales tactics. These features fall broadly under the heading of consumer policy.

It seems hard to define precisely "competition policy" versus "consumer policy". Motta (2004, page 30) suggests that competition policy comprises "the set of policies and laws which ensure that competition in the marketplace is not restricted in such a way as to reduce economic welfare." Whereas according to Vickers (2004) one might define consumer policy in terms of the fundamental problems it seeks to prevent, cure, or remedy, which are: (i) duress and undue sales pressure; (ii) information problems pre-purchase; and (iii) undue surprises post-purchase. Nevertheless, many policies (such as those which act to reduce consumer search costs or switching costs, or which reduce industry advertising costs) could be said to fall under both headings.

For better or worse, there has been a lot more economics informing competition policy than consumer policy. However, in recent years economists have shown a greater interest in consumer policy. In part, this is because the economics profession has recently been energized by behavioral economics, a branch of the discipline which takes into account imperfect consumer decision making—consumers can be less rational, more prone to various biases and temptations.

This paper summarizes for a relatively non-technical audience my own recent work on the economics of consumer protection, which has been done in collaboration with John Vickers and Jidong Zhou.[1] For the most part, and unlike the other papers presented at this conference, this work models consumers as rational agents, and as such it provides rationales for consumer policy which do not need to use recent models of behavioral consumers. In the following sections I present three theoretical models which illustrate the merits and drawbacks of a number of familiar consumer protection policies. First, preventing firms from setting unduly high prices in markets such as credit cards, energy or international mobile telephony may reduce a consumer's incentive to investigate their mar-

[1]For more wide-ranging surveys on the economics of consumer protection, see Vickers (2004) and Armstrong (2008). For an advanced exposition of theoretical models which model consumers as having bounded rationality (and firms as profit-maximizing), see Spiegler (2011).

ket thoroughly. The resulting "model hazard" may well induce firms to *raise* their prices. As such, a safeguard price cap of this form may be a kind of protection which consumers do not need (although it would be welcomed by firms).

Second, policy sometimes aims to prevent firms from rushing their customers' decision making. Sellers may have an incentive to force potential customers to decide then and there whether to buy the product, before the customer has a chance to investigate other—perhaps superior—deals available in the market. When a seller uses this particular sales technique, the result may be a poor match between the consumer and product. In addition, the practice may also lead the seller to set a higher price, which provides another source of consumer harm. While a direct ban on this form of firm behavior may be hard to implement, other common consumer policies such as mandated "cooling off" periods may have the same end result.

Third, in many markets intermediaries act to recommend or otherwise "push" a particular product to consumers. Examples include salesmen giving advice about financial products, doctors giving medical advice, or retailers which make prominent certain products in their shop displays. A common arrangement is for suppliers to pay commission to an intermediary which gives the latter a reward in the event of a sale. A natural worry is that the intermediary then promotes the product which comes with the highest commission, rather than the product which is best for the consumer. In the model discussed, the use of per-sale commission payments acts to raise prices in the market, relative to a market in which consumers pay the intermediary directly for advice, or where uniformed consumers shop randomly.

2 Consumer Protection and Moral Hazard

If consumers are over-protected in their market transactions, "moral hazard" may ensue and they may not pay sufficient attention to making the best choices. As is well understood, if someone is insured, she will take less care protecting her possessions. An efficient insurance contract will trade off the benefits of insurance to risk-averse consumers with the need to ensure that the consumer takes adequate care. Likewise, in markets with complex products or with many suppliers, the consumer needs to invest effort to choose what product is the best for her. For instance, if policy ensures the consumer will face no bad surprises in the small print (of a contract with a bank, for instance), she may be less likely to read the contract at all. As Posner (1969, page 67) put it: "Just as the cheapest way to

reduce the incidence of certain crimes, such as car theft, is by inducing potential victims to take simple precautions (locking car doors), so possibly the incidence of certain frauds could be reduced at least cost to society by insisting that consumers exercise a modicum of care in purchasing, rather than by placing restrictions on sellers' marketing methods."

It seems plausible that consumers learn market skills over time and, moreover, these market skills are often not specific to one market, but spill over to many markets. For instance, the victim of a scam, or an unexpectedly high credit card penalty charge, will usually be more vigilant in future. It does not take many bad experiences with scams to learn the maxim that "if it seems too good to be true, it probably is." Unless a consumer is particularly vulnerable or the product is particularly harmful, it is probably best to let consumers develop their own imperfect rules of thumb to defend themselves in the market. Some consumers will no doubt harm themselves by inexpertly cooking a chicken (say, by not reading the small print of the "cooking instructions"), but the solution is not to remove raw chicken from the market. The general point is that excessive consumer protection may be inimical to the development of market skills in consumers.[2]

To take a specific example, a consumer policy which acts directly to limit price dispersion in such a market could have perverse effects. If price dispersion is reduced, this reduces the incentive for a consumer to become informed, and so is likely to reduce the number of informed consumers. The net result of reduced consumer search could well be that average prices in the market rise rather than fall, thus harming consumers.

Armstrong, Vickers and Zhou (2009) provide formal modelling of this idea.[3] Our market model was an extension of Burdett and Judd (1983), who studied a model where all consumers are rational, and decide whether or not to become better informed about the deals available in the market on the basis of the expected gains from doing so. Armstrong *et al.* assumed the market had a large number of identical firms which supply a homoge-

[2]However, just because there is moral hazard does not mean insurance should not be offered at all. One might balk at permitting sales to the general public of Japanese pufferfish, which is fatal if prepared even slightly incorrectly. A related issue is the widespread use of "use-by" dates on food. Many consumers never use food beyond its use-by date. Given that the use-by date is chosen so that the foodstuff is almost certain to be edible regardless of local conditions (e.g., how often the consumer's fridge is opened), it is plausible that inefficiency arises from this policy. If use-by dates were less widespread (say, in the days when many consumers purchased meat from a butcher rather than a supermarket), consumers would likely have better skills in detecting whether food is edible (e.g., by smell). This is another instance of how arguably excessive protection leads consumers to possess too few market skills.

[3]Much earlier, Fershtman and Fishman (1994) examined the impact of a price cap and showed that the price cap could act to raise expected prices.

neous product to a large number of consumers. For simplicity, the firms' cost of supply is normalized to zero. Consumers are risk-neutral, and all have maximum willingness-to-pay for a unit of the product equal to v. Consumers are endogenously divided into two groups according to their choice of search technology: the *better informed* and the *less informed*. The former observe more prices on average than the latter, but incur a one-off search cost when they choose to become better informed.

In such a market, firms choose their prices randomly and there is price dispersion.[4] In such a market, a consumer who sees more prices will, on average, find a lower price than a consumer observing fewer prices. Suppose that in market equilibrium an informed consumer's expected price is denoted P_I, a less-informed consumer's expected price is $P_U > P_I$, and the fraction of consumers who choose to become informed is λ. Suppose a consumer can choose to use the superior search technology by incurring a (possibly psychological) cost $s \geq 0$. In general, consumers may differ in their cost of acquiring information, and let $s(\lambda)$ be the search cost of the marginal consumer when λ consumers choose to be informed. (The function $s(\cdot)$ is necessarily weakly increasing.) In general, the two expected prices P_I and P_U are decreasing functions of λ (as illustrated for a related model in section 4 below on Figure 1). For a consumer with search cost s to be willing to become more informed, we require that $s \leq P_U - P_I$ so that it is worthwhile to spend s to discover more prices. In equilibrium, consumers will choose to become informed until the final consumer is indifferent. Thus, the fraction λ of consumers who become well informed in equilibrium satisfies

$$P_U(\lambda) - P_I(\lambda) = s(\lambda) \ .$$

To illustrate this discussion, consider an example where the less-informed consumers see just one price, while the more-informed consumers see two prices. If all consumers have search cost $s = v/20$, one can show that approximately 95% of consumers choose to become informed. All consumers make the expected payment (including search costs for those consumers who choose to become informed) of $P_U = P_I + s \approx v/10$.

Consider a policy which aims to protect less informed consumers against unduly high prices. (For instance, a usury law might take this form in a particular credit market, or

[4]More precisely, this requires that there be some of each kind of consumer and that less informed consumers sometimes see just a single price. To understand why firms cannot set predictable prices, suppose to the contrary that each of a given firm's rivals are known to set the price p. Then if p is above marginal cost, the firm can make more profit by slightly undercutting this price, and so selling to all consumers who see its price. If the price p is equal to marginal cost, then the firm can make positive profit by setting its price above cost and selling to those consumers who happen to see only its price.

consumer advocates might suggest such regulation in the energy or telecommunications sectors if some consumers are found to be paying high prices.) That is to say, policy constrains firms to set prices no higher than $\bar{p}$, where $\bar{p} < v$ is the price cap. The imposition of this price cap has pros and cons. If $P_I(\lambda)$ and $P_U(\lambda)$ are expected prices in the absence of regulation, then Armstrong *et al.* show that the expected prices with price cap $\bar{p}$ become respectively

$$\frac{\bar{p}}{v}P_U(\lambda) \text{ and } \frac{\bar{p}}{v}P_I(\lambda) \ .$$

Thus, for given λ, the intervention benefits both the informed and the uninformed consumers since the prices they pay are proportional to $\bar{p}$. But the incentive to become informed, i.e., the gap between the two expected prices, is also proportional to $\bar{p}$ for given λ, and so the policy induces the number of informed consumers to fall.

Consider imposing the price cap $\bar{p} = v/2$ in the above numerical example, so that maximum allowed prices are halved. In this case the fraction of informed consumers falls to $\lambda \approx 0.74$, so that the number of uninformed consumers rises about 5-fold as a result of the policy. Each consumer pays $(\bar{p}/v)P_U$, which is now increased by about 70% to $0.17 \times v$. Industry profit more than *doubles* as a result of the imposition of the price cap. Thus, the "perverse" effect of this particular consumer policy is substantial in this example.

Beyond this numerical example, when does imposing a price cap harm consumers? In the special case where all consumers have the same search cost s, provided the price cap is not so tight that all consumers cease searching, the imposition of a price cap is sure to make all consumers pay higher expected prices. Thus, the numerical example is not a fluke, and is rather a robust phenomenon. Although the direct effect of a price cap is to reduce prices, the indirect effect of reduced search lessens each firm's demand elasticity so much that prices on average go up. This formalises a claim sometimes made informally, which is that imposing price controls on an oligopoly market could raise equilibrium prices. One intuition for such a claim is that a price cap acts as a focal point for tacit collusion. For instance, Knittel and Stango (2003) examine the credit card market in the United States in the period 1979-89. There, usury laws in many states put a ceiling (often of 18%) on the interest rates which credit cards could levy. Knittel and Stango (2003, Table 3) show how, for much of this period, average interest rates were higher in those states with a ceiling than in those states without any controls. They interpret this observation as evidence that price caps can encourage tacit collusion via the policy-induced focal point. The search-theoretic model in Armstrong *et al.*, however, provides another way to interpret this data. In our model, pricing is entirely non-cooperative, and tacit collusion plays no role. Rather, price

controls soften competition by blunting consumers' incentives to search for good deals.

If consumers differ in their costs of acquiring market information, imposing a price cap causes fewer consumers to cease becoming informed. If the search cost curve $s(\lambda)$ is sufficiently steep, a price cap will then benefit consumers. Consider for instance the limiting case where an *exogenous* fraction of consumers λ are informed while the remaining consumers are uninformed. This situation could be interpreted as there being a fraction λ of consumers have zero search cost and the remainder have an infinite search cost; or we could take a behavioural interpretation, that a fraction $1-\lambda$ of consumers are "naive" and mistakenly believe there is no benefit to shopping around. (This model is essentially Varian's (1980) model of sales.) When λ is constant, the imposition of a price cap is unambiguously beneficial for both groups of consumers (since their expected prices fall), and harms industry profits. Thus, we can conclude that the impact of a price cap on consumer welfare depends in this model on the fine details of the distribution of search costs in the population of consumers.

It would be useful in future work to extend this stylized model to richer settings. For instance, it is not particularly common to impose caps on headline prices in oligopoly markets. Rather, price controls might be applied to "small print" charges in a contract, or minimum quality standards might be imposed on aspects of product quality. It would be worthwhile to extend this model so that consumers must expend effort to understand these less salient aspects of a firm's offer. For instance, could the introduction of a minimum quality standard sometimes lead to lower average quality in the market?

Armstrong *et al.* also consider an alternative setting in which consumers have the ability to "opt out" of intrusive marketing. A popular consumer protection policy is to introduce a "do not call" list, and when someone signs up to such a list marketers are not permitted to make cold-calls to this person. Again, this policy sounds beneficial to consumers, as this form of marketing can be irritating. However, to the extent that this form of marketing allows recipients to become more informed about deals available in the market (albeit at the "search cost" of having to endure the marketing efforts), such a policy again has pros and cons. Prices are pushed downwards when a greater proportion of consumers are well informed, and so when many consumers choose to opt out of marketing this impacts negatively on prices. The net result can be that consumers are harmed when the "do not call" list is introduced. Indeed, firms may welcome this particular consumer policy, as it relaxes price competition in their markets. (For the same reason, historically firms have often supported measures which restrict price advertising.)

3 Rushed Decision Making

One controversial sales method forces the consumer to decide *quickly* whether to buy. Methods of encouraging a quick decision include a seller refusing to sell to a customer unless she buys immediately, or in less extreme cases the seller tells the potential customer that she will pay a higher price if she decides to purchase at a later date. In his account of sales practices, Cialdini (2001, page 208) reports:

> Customers are often told that unless they make an immediate decision to buy, they will have to purchase the item at a higher price later or they will be unable to purchase it at all. A prospective health-club member or automobile buyer might learn that the deal offered by the salesperson is good for that one time only; should the customer leave the premises the deal is off. One large child-portrait photography company urges parents to buy as many poses and copies as they can afford because "stocking limitations force us to burn the unsold pictures of your children within 24 hours". A door-to-door magazine solicitor might say that salespeople are in the customer's area for just a day; after that, they, and the customer's chance to buy their magazine package, will be long gone. A home vacuum cleaner operation I infiltrated instructed its sales trainees to claim that, "I have so many other people to see that I have the time to visit a family only once. It's company policy that even if you decide later that you want this machine, I can't come back and sell it to you."

A related example is the practice in some academic disciplines for journals to make exploding offers to authors, requiring them to commit to publish with them before they find out whether other, perhaps better, journals are willing to publish their article. Because of the inefficient decision-making the use of exploding offers induces, recently a number of law journals have agreed to cease their practice of making exploding offers to authors.[5]

A less extreme sales tactic is to offer a discount for immediate sale. A home improvement company might offer its potential customers a regular price for the agreed service, together with a discounted price if the customer agrees immediately. Similarly, a prospective tenant might be offered an apartment for $900 per month but to whom the landlord offers $850 if she agrees immediately, or a car dealer tries to close a deal who offers a further $500 off the price if the buyer accepts now, so (as he claims) he can then make his sales quota for that month.

[5] See the letter published online at www.harvardlawreview.org/Joint-Letter.pdf.

Inducements to make a quick decision can limit a consumer's ability to make a well-informed decision, which in turn can harm market performance. Public policy has attempted to address this problem. For instance, the *Unfair Commercial Practices Directive,* adopted in 2005 across the European Union, prohibits in all circumstances "Falsely stating that a product will only be available for a very limited time, or that it will only be available on particular terms for a very limited time, in order to elicit an immediate decision and deprive consumers of sufficient opportunity or time to make an informed choice." However, in practice the enforcement of such laws is often difficult. A more efficient method to tackle the issue may sometimes involve less direct means. For example, exploding offers could in essence be prohibited by mandating a "cooling off period", so that consumers have the right to return a product within some specified time after agreeing to purchase. (They could then return a product if they subsequently find a preferred option.) Many jurisdictions impose cooling off periods for some products, especially those sold in the home.

Armstrong and Zhou (2011b) provide a formal model to examine a seller's incentive to encourage rushed decisions, by discriminating against those customers who wish to buy its products later. It is natural to study this issue in the context of sequential search, where consumers search for a suitable product and/or for a low price.[6] Of course, the sales tactic only works in those situations where sellers can distinguish new visitors from people who have returned to buy only after the initial sales pitch. In the majority of markets this is not possible. (A supermarket, for instance, keeps no track of a consumer's entry and exit from the store.) Nevertheless, in many markets such discrimination is feasible. A sales assistant may tell from a potential customer's questions or demeanor whether she has paid a previous visit or not, or may simply recognize her face. In online markets, a retailer using tracking software may be able to tell if a visitor using the same computer has visited the site before. Sometimes—as with job offers, automobile sales, tailored financial products, medical insurance, doorstep sales, or home improvements—a consumer needs to interact with a seller to discuss specific requirements, and this process reveals the consumer's identity.

[6] As mentioned in the introduction, we used a model with rational consumers. There are many other methods to induce sales which rely on more psychological factors. These include attempts to make the prospective buyer "like" the seller (e.g., by claiming similar interests, family or social background) or attempts to make the buyer feel obligated to the seller (e.g., by means of a "free gift"). Cialdini describes these and other sales techniques in more detail. However, it is often unclear what role consumer policy has to play in combatting these kinds of sales tactics.

In such situations, there are two reasons why a firm might wish to discriminate against those consumers who buy later. First, there is a *strategic* reason, which is to deter a potential consumer from going on to investigate rival offers. If a consumer cannot return to a seller once she leaves, this increases the opportunity cost of onward search, as the consumer then has fewer options remaining relative to the situation in which return is costless. Second, the observation that a consumer has come back to a seller after sampling other options reveals relevant *information* about a consumer's tastes or the prices she has been offered elsewhere, and this may provide a profitable basis for price discrimination. A seller may charge a higher price to those consumers who have already investigated other sellers, because their decision to return indicates they are unsatisfied with rival products. The former motive is most relevant when firms announce their buy-later policies in advance (and stick to their policies), while the latter is more important when firms have less commitment power.

A simple framework to think about these issues is the following. A single seller supplies a product which yields gross utility u to a consumer, where u varies across consumers such that the fraction of consumers with $u \geq p$ is described by the demand function $Q(p)$. The key twist to the model is that the consumer's outside option (her utility received if she does not buy the seller's product), denoted by v, is a random variable which the consumer does not know until she leaves the seller. (The parameter v might represent the uncertain value of other deals available from alternative sellers, for instance.) If the seller chooses price p, the consumer's net surplus from the seller is $u-p$. If the seller allows the consumer to investigate her outside option before deciding whether to buy, the consumer will always wait to discover the outside option (in case v happens to be large), and then return to buy whenever $u - p \geq v$. With this method of selling, the probability that the consumer buys its product is the expected value of demand, denoted $\mathbb{E}Q(p + v)$ (where the expectation takes place with respect to the outside option v). If instead the seller forced the consumer to decide to buy *before* she can find out v, with price p the consumer with gross utility u will accept this exploding offer whenever $u-p \geq \mathbb{E}v$, where $\mathbb{E}v$ is the expected value of the outside option. (Here, we assume the consumer is risk-neutral.) Thus, the probability of a sale with this high-pressure sales technique is $Q(p+\mathbb{E}v)$. According to Jensen's Inequality, $\mathbb{E}Q(p+v)$ is smaller than $Q(p+\mathbb{E}v)$ if the demand curve $Q(\cdot)$ is concave over the relevant range, and it is greater if the demand curve is convex.

Thus, in this simple setting, the incentive to make an exploding offer depends on the shape of the seller's demand curve: with a concave demand curve the seller has an incentive

to use this form of high-pressure selling.[7] The basic trade-off involved is as follows. When the seller makes an exploding offer, this makes the consumer more likely to accept the offer immediately if she likes it, but it prevents her, in the event that she has only a moderate payoff from the offer, from coming back if she discovers her outside options is even worse. When the demand curve is concave, the first effect dominates. For a given price p, the consumer is harmed when the seller makes an exploding offer, since she obtains her ideal outcome when free recall is allowed while an exploding offer leads to an inefficient outcome for many realizations of (u, v). In addition, the use of an exploding offer may induce the seller to alter its chosen price; it will raise the price when an exploding offer is made if the demand function $Q(p + \mathbb{E}v)$ is less elastic than the demand function $\mathbb{E}Q(p + v)$. In general, this comparison is ambiguous, and depends on the concavity or convexity of the *slope* of demand. However, the typical pattern seems to be that the seller raises its price when it makes an exploding offer. In such cases, the use of exploding offers has a double disadvantage: the tactic induces a poor match between consumers and products and it raises the price consumer must pay.

While firms have an incentive to make an exploding offer in the relatively restrictive case where the demand curve is concave, they have an incentive to offer a buy-now discount much more widely. Indeed, Armstrong and Zhou (2011b) show that a firm has such an incentive under the mild condition that the demand curve is log-concave. Although the sales tactic is framed as a discount (e.g., "buy my product now and you'll save 10% off my usual price"), it turns out that when a firm engages in this form of price discrimination *both* its prices often rise relative to a situation where the firm offers a uniform price to its customers. Again, in such cases the sales tactic induces a poor product match and higher prices.

An alternative method of discriminating against prospective buyers who leave and then return is to implement an *unannounced* price hike. When searching for air-tickets online, a consumer may find a quote on one website, go on to investigate a rival seller, only to return

[7] This result continues to hold even if the seller cannot commit not to serve a returning customer, provided some consumers are "credulous" and believe the sales patter. In reality, a doorstep seller, say, may be only too pleased to return to sell if a customer calls to say she does in fact want the item. In such cases, the exploding offer is not a credible sales strategy. However, if some consumers do anyway believe the salesman's claim that he is "in the area that day only" and the purchase decision must be immediate, the salesman has an incentive to claim to make an exploding in order to influence the decision of these credulous consumers. (The sophisticated consumers are not taken in or otherwise affected by the salesman's claims.)

to the original website to find the price has mysteriously risen. Or a consulting firm may be approached by a company wanting antitrust advice and a fee is chosen, but if the company returns some weeks later after trying rival consultants (who are too expensive, or perhaps turn out to be conflicted), it may find the fee has increased. To analyze such cases, we relaxed the assumption that firms commit to their buy-later price when consumers make their first visit. Then it is often the case that the seller does wish to raise its initial price when a consumer comes back to buy later.

For instance, suppose a consumer incurs a (possibly small) intrinsic cost $r > 0$ in order to return to the seller after investigating the outside option. If the seller initially offers the price p and the consumer anticipates that this price will remain valid if she comes back to buy later, then any consumer who buys later must have preferences such that $u - p - r \geq v$. (If she is willing to come back, then her surplus at the seller, $u - p$, must exceed the outside option v by enough to compensate her cost of returning.) Therefore, the seller can raise its price from p to $p + r$ and not induce any of these returning consumers to be driven back to the outside option. In fact, a similar argument shows that there can then be *no* buy-later price which is accurately anticipated by consumers. It follows that the only credible outcome when firms have no commitment power at all is that the seller makes an exploding offer and the return market collapses. An inability to commit to its buy-later policy will therefore amplify a firm's incentive to discriminate against those consumers who buy later.

4 Commission-based Selling

As discussed in the two previous sections, consumers are often initially imperfectly informed about the deals available, and must invest effort to find out where to obtain a reasonable product at a reasonable price. A consumer may sift sequentially through the options available until she finds one which is satisfactory (rather than the *best* available in the market). In such a market, a seller has an advantage if it is encountered early on in a consumer's search process. In a few situations it makes sense to suppose that consumers search randomly through available options, in which case no firm is privileged relative to its rivals. In many circumstances, however, consumers consider options in a non-random manner, and choose first to investigate those sellers or products which have high brand recognition, which are known to have a low price, which the consumer has purchased previously, which are recommended by an intermediary, or which are prominently displayed within a retail environment.

Armstrong and Zhou (2011a, section 1.1) consider a setting where firms market their products by offering financial inducements to intermediaries. The formal model assumed that sellers could not observe—or contract on—the products the intermediary chooses to promote, and to give an incentive to promote its product a seller pays a per-sale commission fee to the intermediary. This sales method is often used in one-to-one sales environments such as for financial services. In this model, the intermediary chooses to "recommend" the product which pays the highest commission, and uninformed consumers are steered towards the more expensive product. This could be construed as a form of mis-selling. Because sellers compete to become prominent by offering high commissions, this pushes up a seller's marginal cost of supply, and so equilibrium retail prices are high relative to a market with random consumer search.

We studied a variant of Varian (1980) in which his framework is modified to allow a single intermediary (or "salesman" for brevity in the following) to steer the uninformed portion of consumers towards a particular product. In more detail, a number of symmetric sellers costlessly supply a homogenous product (life insurance, say) which all consumers value at v. We assume that this product *must* be sold via the salesman. An exogenous fraction λ of well-informed consumers costlessly observe the two retail prices, and buy from the cheapest supplier. The remaining fraction $1-\lambda$ of consumers will only consider a single product and buy that product if its price is below v. (These consumers may have very high search costs, or are susceptible to the marketing efforts of the salesman and follow his recommendation.) Hence, the salesman has the ability to steer these $1-\lambda$ uninformed (or "credulous") consumers to buy any particular product. Suppose that a firm chooses its retail price, p, and commission rate, b, simultaneously (and simultaneously with its rivals). This firm pays commission b to the salesman every time a sale of its product is made. We assume that the salesman cannot levy charges on consumers, and so aims to maximize his income from commission payments.

In this setting it is clear that the salesman will choose to promote the high-commission product, regardless of how the two retail prices compare (as long as prices do not exceed v). This is because the salesman's marketing effort cannot influence the choice of the informed consumers at all, but fully determines the choice made by the uninformed consumers. Hence, the salesman will direct the uninformed consumers towards the product which pays a higher commission rate. It is also clear, as in the work described in section 2, that sellers choose their retail prices and commission payments randomly. In the 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. Moreover, this incentive also increases with the proportion of uninformed consumers. Since high commissions are associated with high retail prices, the salesman promotes the highly priced product due to the high commission he then receives. This is a form of mis-selling, since uninformed (or credulous) consumers are directed towards the more expensive product.

There are two natural benchmarks with which to compare the outcome when commissions are paid. The first benchmark is when there is no salesman, and the uninformed consumers buy randomly from one of the firms. In this case the framework reduces exactly to Varian (1980)'s model. We show that retail prices are higher when firms pay commissions to a salesman to promote their product relative to the situation with random search. This is due to the competition between firms to offer high sales commissions to have their product promoted, which artificially inflates the marginal cost of selling a product. However, whether firms enjoy greater profits when they pay commission is ambiguous. In the case of two suppliers, without commission payments each firm makes expected profit $\frac{1}{2}(1-\lambda)v$, while in the regime with commissions a firm makes expected profit $\lambda(1-\lambda)v$. Thus, more profit is obtained with commission payments when $\lambda > \frac{1}{2}$, so that the uninformed consumers are in the minority. But when the uninformed consumers are in the majority, the two firms end up playing a prisoner's dilemma due to the fierce competition to become prominent.

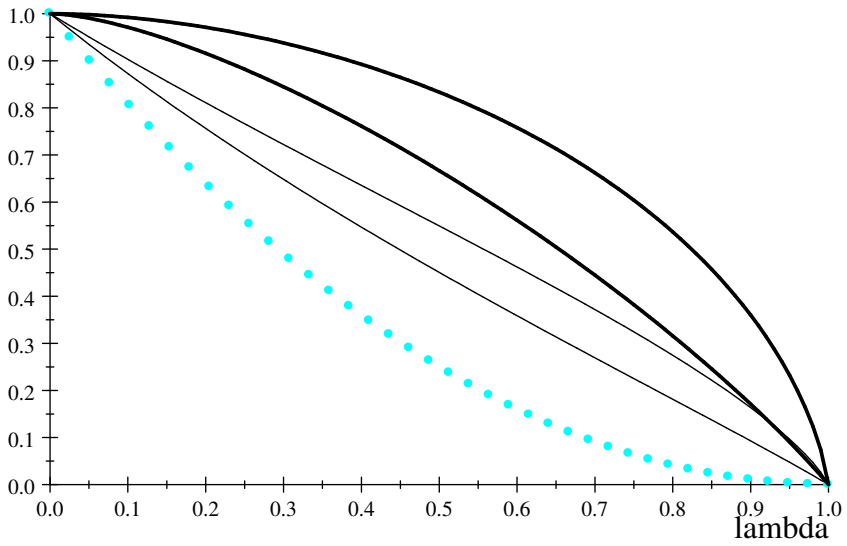

Figure 1: *Expected prices and commissions in three regimes*

Figure 1 plots the expected prices paid in these two regimes as a function of λ, the

proportion of informed consumers. (Here, $v = 1$.) The two bold lines depict expected prices when commissions are paid, where the upper of these lines is the expected price paid by the uninformed consumers and the lower line is the expected price paid by the informed consumers. The dotted line represents the expected commission paid to the salesman. The two feint lines depict the corresponding prices in Varian's model where no commissions are paid and search is random. The two regimes have the same outcome for consumers when $\lambda = 0$ (when the monopoly price $p = v$ is chosen for sure) and when $\lambda = 1$ (when the competitive price $p = 0$ is chosen). However, for intermediate values of λ, the prices paid in the commission regime are substantially higher than when no commissions are paid. Indeed, in most cases an uninformed consumer in the no-salesman regime pays a lower price than even the informed consumers do in the commission regime.

The second benchmark with which to compare the outcome with commission payments is to suppose that the salesman is necessary for consumers to buy the product (unlike the benchmark with random search), but now the salesman is paid by consumers rather than by sellers.[8] Suppose that when the salesman is paid by consumers, say in the form of a lump-sum consultation fee, he directs the uninformed consumers to the cheaper product. (This might be because, all else equal, he has a small intrinsic preference for recommending the appropriate product to consumers.) In this case, all consumers buy the cheaper product and in Bertrand fashion the sellers are forced to set retail prices equal to cost. Thus, suppliers are harmed when this policy is introduced, relative to both the commission regime and the random search regime. The outcome for consumers depends on how much they have to pay the salesman for his advice. One assumption is that the consultation fee is set equal to the revenue the salesman received under the commission regime, so that the salesman is indifferent between the two regimes. (Perhaps the advice industry needs to be supportive of a policy shift from a commission-based model to a consumer-fee model.) In this case, the expected total price—the price for the product plus the fee to the salesman—paid by any consumer is simply the dotted line on Figure 1. From the figure it follows that all consumers are better off when they pay the salesman compared to when suppliers pay the salesman. In fact, they are also better off when they pay the salesman than when they search randomly (where prices are the feint lines on the figure).

This section has described a model where firms attempt to influence a salesman's mar-

[8]The UK regulator, the *Financial Services Authority*, published rules in March 2010 concerning how financial advice can be remunerated. The rules state that an advisor will not be able to accept commission for recommending products, and the consumer fee for advice must be agreed between the consumer and the advisor, rather than between the seller and the advisor.

keting efforts by means of per-sale commission payments. The salesman gives prominence to the product which pays the highest commission, and in equilibrium this entails steering uninformed consumers towards the more expensive product. Competition between sellers to set the highest commission means that the marginal cost of supply is inflated and equilibrium retail prices are high. Therefore, the outcome for consumers, both informed and uninformed, is poor: worse than the situation without commission payments where the uninformed shop randomly, and far worse than the situation in which consumers pay directly for advice. This model therefore gives some support to consumer policies which restrict the use of commission payments as a marketing tactic.[9]

This discussion considered an environment in which sellers could not observe the marketing efforts of the intermediary, and so induced effort from the intermediary with the use of per-sale commissions. A by-product of this arrangement is that a seller's marginal cost of supply is artificially inflated, and consumers are harmed by high retail prices which result. In other environments, sellers can observe the intermediary's marketing strategy, and so there is no need to give incentives *ex post* for the intermediary to promote the product. (For instance, a publisher can observe whether a bookstore does in fact promote its book as the "book of the month".) As such, it is then often more natural to suppose that payments for promotion are lump-sum rather than per-sale, with the result that retail prices are not necessarily adversely affected.[10] Indeed, as discussed in Armstrong and Zhou (2011a, section 1.2), lump-sum payments for product promotion may actually be welfare-enhancing, as sellers with better (or cheaper) products may well be prepared to pay the most for being promoted in this way, and so consumer will end up being guided in the appropriate direction.

[9] Inderst and Ottanviani (2011) present an alternative model of potential 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 sellers set high enough commissions.

[10] One UK bookstore was alleged in 2006 to charge publishers £50,000 a week to guarantee a book "a prominent position in the store's 542 high street shops and inclusion in catalogues and other advertising". A trade body suggested that 70 per cent of publisher promotional budgets were spent on so-called "below-the-line" schemes operated by bookshops rather than more traditional advertising. For more details, see the article in the (UK) *Sunday Times* by Robert Winnett and Holly Watt titled "£50,000 to get a book on recommended list", 28 May 2006.

5 Concluding Comments

I have presented three theoretical models which aim to shed light on the pros and cons of a number of common consumer protection policies. We saw that some support could be given to policies which seek to prevent rushed decision making and which seek to control the use of per-sale commission payments as a method of giving incentives to sales intermediaries. In both of these cases, it was not clear from the models whether "more competition" would be another way to solve these problems. We also saw how an apparently pro-consumer policy which limits maximum prices in the market might backfire, and lead firms to raise their average prices.

These policies fall under the headings of combatting either (i) sales pressure or (ii) information problems before purchase (using the taxonomy in Vickers (2004)). In future work it will important to understand better the problems which emerge with (iii) surprises after purchase. For instance, in what circumstances should regulation control terms in the "small print" of consumer contracts, and if so, how should it do so? For example, many consumers are known to overlook contractual terms such as unauthorized overdraft charges levied by banks, or call charges levied by mobile telephone networks when a subscriber makes more calls than their allowance. Is there a role for consumer policy to control these charges, and if so, at what level? To what extent is the moral hazard problem analyzed in section 2 likely to re-emerge with small print regulation?

References

Mark Armstrong (2008), "Interactions between competition and consumer policy", *Competition Policy International* 4(1): 97-147.

Mark Armstrong, John Vickers and Jidong Zhou (2009), "Consumer protection and the incentive to become informed", *Journal of European Economic Association* 7(2-3): 399-410.

Mark Armstrong and Jidong Zhou (2011a), "Paying for prominence", *Economic Journal* 121: F363-395.

Mark Armstrong and Jidong Zhou (2011b), "Exploding offers and buy-now discounts", mimeo.

Kenneth Burdett and Kenneth Judd (1983), "Equilibrium price dispersion", *Econometrica* 51: 955-969.

Robert Cialdini (2001), *Influence: Science and Practice*, Allyn and Bacon (fourth edition).

Chaim Fershtman and Arthur Fishman (1994), "The 'perverse' effects of wage and price controls in search markets", *European Economic Review* 38: 1099-1112.

Roman Inderst and Marco Ottaviani (2011), "Competition through commissions and kickbacks", *American Economic Review* (forthcoming).

Christopher Knittel and Victor Stango (2003), "Price ceilings as focal points for tacit collusion: evidence from credit cards", *American Economic Review* 93: 1703–1729.

Massimo Motta (2004), *Competition Policy: Theory and Practice*, Cambridge University Press.

Timothy Muris (2002), "The interface of competition and consumer protection", paper presented at Fordham Corporate Law Institute's 29th Annual Conference on International Antitrust Law and Policy, New York.

Richard Posner (1969), "The Federal Trade Commission", *University of Chicago Law Review* 37(1): 47-89.

Ran Spiegler (2011), *Bounded Rationality and Industrial Organization*, Oxford University Press.

Hal Varian (1980), "A model of sales", *American Economic Review* 70(4): 651-569.

John Vickers (2004), "Economics for consumer policy", *Proceedings of the British Academy* 125: 287-310.