JENA ECONOMIC RESEARCH PAPERS

2009 – 084

Spurious complexity and common standards in markets for consumer goods

by

**Alexia Gaudeul
Robert Sugden**

www.jenecon.de

ISSN 1864-7057

The JENA ECONOMIC RESEARCH PAPERS is a joint publication of the Friedrich Schiller University and the Max Planck Institute of Economics, Jena, Germany. For editorial correspondence please contact markus.pasche@uni-jena.de.

Impressum:

Friedrich Schiller University Jena
Carl-Zeiss-Str. 3
D-07743 Jena
www.uni-jena.de

Max Planck Institute of Economics
Kahlaische Str. 10
D-07745 Jena
www.econ.mpg.de

© by the author.

Spurious complexity and common standards in markets for consumer goods*

Alexia Gaudeul and Robert Sugden

Jena Graduate School, Friedrich Schiller University Jena

Centre for Competition Policy and School of Economics, University of East Anglia

14 October 2009

Abstract: It has been argued that cognitively constrained consumers respond sub-optimally to complex decision problems, and that firms can exploit these limitations by introducing spurious complexity into tariff structures, weakening price competition. We model a countervailing force. Restricting one's choices to the most easily comparable options is a psychologically well-attested heuristic. Consumers who use this heuristic favour firms that follow common conventions about tariff structures. Because a 'common standard' promotes price competition, a firm's use of it signals that it offers value for money, validating the heuristic. This allows an equilibrium in which firms use common standards and set competitive prices. (100 words)

JEL classifications: D83 (search, learning, information and knowledge); L13 (oligopoly and other forms of imperfect competition); L51 (economics of regulation).

Keywords: common standard, spurious complexity, cognitive limitations.

* Earlier versions of this paper were presented at a conference on consumer behaviour and bounded rationality at the ESRC Centre for Economic Learning and Social Evolution, University College London, at the 35th Conference of the European Association for Research in Industrial Economics in Toulouse, and at various workshops and seminars. We thank participants in these meetings, Todor Tsankov, Nick Watson, an editor and an anonymous referee for comments and suggestions. Sugden's work was supported by the Economic and Social Research Council of the UK (award no. RES 051 27 0146).

Addresses

Alexia Gaudeul, Jena Graduate School (Human Behaviour in Social and Economic Change), Friedrich Schiller University Jena, Bachstraße 18k, 07743 Jena, Germany; a.gaudeul@gmail.com

Robert Sugden, Centre for Competition Policy and School of Economics, University of East Anglia, Norwich NR4 7TJ, UK; r.sugden@uea.ac.uk

There is growing evidence that consumers can find it difficult to process complex decision problems. As a result, they may fail to choose in accordance with what, after sufficient reflection, they would acknowledge to be their own best interests. The recognition of this problem by behavioural economists is producing a literature which advocates regulatory interventions to simplify consumers' choice problems (Sunstein and Thaler, 2003a, 2003b; Camerer et al, 2003; Thaler and Sunstein, 2008). A complementary literature in industrial organisation is investigating whether profit-maximising firms can exploit consumers' cognitive limitations by introducing spurious complexity into tariff structures. The typical finding is that firms have incentives to follow such strategies, and that their doing so tends to raise prices for cognitively or informationally constrained consumers (Ellison, 2005; Gabaix and Laibson, 2006; Spiegler, 2006; Ellison and Ellison, 2009). These findings appear to strengthen the case for regulation by showing that, in the absence of regulation, consumers do not merely have to navigate the 'natural' complexity of competitive markets; they also have to cope with unnecessary complexity which has been deliberately created to confuse them.

In this paper, we argue that these literatures neglect an important countervailing force: the *common standard effect*. The essential idea is that consumers' choice problems are made less complex if competing firms follow common conventions about tariff structures, package sizes, labelling, and so on. Cognitively constrained consumers can be expected to simplify decision problems by focusing attention on those options that are most easily comparable, and hence to favour products which meet common standards. By facilitating comparisons between products, common standards promote competition between the firms that follow them. But, precisely because they promote competition, they also signal that goods that meet common standards are likely to offer good value for money; this reinforces the tendency for consumers to favour common standards. If consumers behave in this way, profit-seeking firms are penalised for deviating from common standards.

We present a model of a market in which, in the absence of common standards, consumers would find it difficult to make accurate comparisons between the tariffs of competing firms, allowing firms to set prices above the competitive level. We investigate

the conditions under which, despite this opportunity for the exploitation of consumers' cognitive limitations, common standards can be self-sustaining.

The common standard effect can be distinguished from other market mechanisms which promote the simplification of consumers' choice problems. In particular, it should be distinguished from those mechanisms which work through the incentive for individual firms to build reputations as trading partners who provide value for money, rather than seeking to trap unwary consumers. The common standard effect is a complementary but distinct mechanism, which works at the level of the market rather than the firm. Common standards are market-wide conventions. Firms reveal themselves as offering value for money, not by signalling their *individual* identities as reliable trading partners, but by displaying features that are characteristic of reliable firms *in general*.

The hypothesis that price competition can be frustrated by the bounded rationality of economic agents must be distinguished from the more general idea that deviations from traditional rationality assumptions have implications for the behaviour of profit-maximising firms. Among the non-standard preferences or 'biases' whose effects have been investigated are dynamic inconsistency (DellaVigna and Malmendier, 2004, 2006; Eliaz and Spiegel, 2006), over-confidence (Grubb, 2008), and loss aversion (Heidhues and Koszegi, 2004). This distinction is significant in the light of Sugden's (2004, 2009) demonstration that competitive markets can be efficient in generating opportunities for individual consumers, whether or not they have well-defined or consistent preferences, provided that they buy only at the lowest price posted in the market and sell only at the highest. Thus, if the normative criterion is one of opportunity, competitive markets have desirable properties independently of the rationality or irrationality of consumers' preferences. But if price competition itself were obstructed by the bounded rationality of consumers, the opportunity-based defence of markets would be compromised.

1. Common standards

A common theme in the discussion of spurious complexity is that price competition can be blunted if the prices that are posted by sellers provide the information that is relevant for price comparisons only indirectly, and if that information can be reconstructed only at some cost in terms of search or cognitive effort, or with some probability of error. For example, a package of services that would more naturally be priced as a single entity can be broken

down into separately-priced components. This imposes cognitive costs on buyers who want to know the price of the whole package, and introduces the possibility (which sellers may find ways of making more probable) that buyers overlook some components in calculating that price. This form of complexity is exemplified by some forms of add-on pricing (Ellison, 2005; Gabaix and Laibson, 2006; Ellison and Ellison, 2009) and, for contingency-related services such as health insurance or retail banking, the practice of posting different prices or reimbursement schedules for different contingencies (Spiegler, 2006). Another way in which posted prices may not allow direct comparisons across products is through the use of different quantity units. For example, different sellers of a product may use different package sizes and display per-package rather than per-unit prices, or use different conventions of measurement, such as different statistics for expressing interest rates.

Common to all these examples is the idea that price comparisons are facilitated if, in providing price information, competing suppliers use a common standard. By ‘standard’, we mean a well-defined statistic (or array of statistics) which, if known by consumers, is informative about the pricing schedule of each of a set of competing firms and allows comparisons between those schedules. If two or more firms express their prices using the same standard, that standard is ‘common’ to them. In contrast, a standard which is used by only one firm is ‘individuated’. If suppliers use individuated standards when an informative common standard could have been used instead, spurious complexity is created.

The imposition of common standards is a familiar form of market regulation. For example, there are EU regulations which require the display of prices per metric unit of weight or volume, which require prices to be expressed inclusive of taxes, and which require interest charges to be expressed using a standard annual percentage rate formula. Until 2009, EU regulations also required many pre-packaged consumer products to be sold only in limited ranges of package sizes so as to facilitate price comparisons. In Singapore, the mark-up that hotels charge on international phone calls is capped by law (Gabaix and Laibson, 2006: 531).

However, it is illuminating to consider whether common standards can be sustained in the absence of regulation. We suggest that some common standards *have* emerged spontaneously and have persisted for relatively long periods as self-enforcing conventions. For example, consider the practice of add-on pricing. Writers on this topic can point to

familiar examples of high-priced add-ons such as phone calls, internet connections and minibar drinks in hotels. But it is equally instructive to notice that generic references to such examples are possible only because readers can be expected to know that, across a whole market sector, particular services are *standardly* (over-)priced as add-ons, while others are *standardly* included in the basic price. Thus, most mid-range hotels treat minibar drinks as add-ons but do not charge for the use of in-room coffee-making facilities, toiletries or trouser-presses. Similarly, the prices posted by ‘full-service’ airlines have for many years standardly included normal check-in baggage and on-board meals and soft drinks.¹ We conjecture that such conventions are self-enforcing because firms expect that unilateral deviations would be perceived by consumers as sharp practice, and would induce a loss of confidence in the meaningfulness of their posted prices.

Of course, there are now expanding sectors of budget hotels and budget airlines which do not follow full-service conventions. But it is significant that firms in these budget sectors typically advertise themselves as offering new business models, rather than trying to pass themselves off as full-service suppliers. For example, the website of Travelodge (the leading UK budget hotel chain) declares ‘We keeps costs low by cutting out things you don’t need’ and ‘Most people staying in a hotel simply want a clean, comfortable place to get a good night’s sleep, and are happy to forgo the unnecessary “frills” offered in other stuffy, over-priced establishments’. Rather than providing these frills as high-priced add-ons, Travelodge either does not supply them at all (there are no bath mats), offers them free of charge to customers who specifically ask for them (extra towels are available at reception), or sells them at low prices (there are no chocolates on your pillow, but snacks can be bought at a vending machine).² It seems that Travelodge is deliberately trying to educate consumers in a new set of conventions about hotel services and pricing.

An example of the emergence of a standard was observed by Ellison and Ellison (2009: 434–435) during their study of Pricewatch, an internet search engine for the purchase of computer parts. At first, Pricewatch did not collect information on shipping costs, and ‘it was not uncommon for firms to list a price of \$1 for a memory module and inform consumers of a \$40 “shipping and handling” fee at check out’. To protect the value of the price-comparison service that *it* was supplying, Pricewatch required all firms to offer a standard form of ground shipping with prescribed maximum charges. It also added a column to its listing in which firms could display shipping cost; customers were explicitly warned to be wary of firms which left the column blank. As a result, many (but not all)

firms chose to charge and display the prescribed maximum. This is a ‘common standard’ in our sense. For a typical Pricewatch user, the package for which price information is relevant is the combination of a physical product and shipping; a convention by which all suppliers post the same shipping fee reduces the dimensionality of price comparisons. In this case, the emergence of a common standard seems to have been partly the result of ‘regulation’ by a market intermediary and partly spontaneous.

Intuitively, it seems that we are observing a balance of forces, some of which support the emergence of common standards while others support deviation from those standards. The existing literature on spurious complexity has concentrated on the latter. Our paper is an attempt to redress the balance.

2. The asymmetric dominance and shortlisting heuristics

Imagine that EU and UK regulations on the display of prices have been repealed and many British firms have gone back to using Imperial units. You have just driven off a ferry at Dover and need to fill up your car with petrol. You see three filling stations, A, B and C, with no obvious differences between them apart from the prices they are displaying. These are £1.05 per litre (at A), £4.75 per Imperial gallon (at B), and £4.76 per Imperial gallon (at C). You are not good at arithmetic, you have no pocket calculator, and you can’t remember the conversion factor between gallons and litres.³ Which filling station should you choose?

Clearly you should not choose C, since B is certainly cheaper than C. One simple heuristic is to choose B on the grounds that B is at least cheaper than *some* alternative, while A is not.

This way of reaching a decision can be interpreted as an instance of the more general heuristic of *asymmetric dominance*. Consider any three options x , y and z that can be located on two dimensions of value; y dominates z while x neither dominates nor is dominated by either y or z . (For example, suppose that x , y and z are different brands of AA battery and the two dimensions are cheapness and durability; the brands are ranked $\langle x, y, z \rangle$ in descending order of cheapness and $\langle y, z, x \rangle$ in descending order of durability.) For a wide class of decision problems, it has been found that individuals are more likely to choose y when the opportunity set is $\{x, y, z\}$ than when it is $\{x, y\}$. This effect has been found both for human and non-human decision-makers; it is recognised (as the *decoy effect*) in the literature of marketing (Huber, Payne and Puto, 1982; Shafir, Simonson and Tversky,

1993; Shafir, Waite and Smith, 2002). It seems that individuals who have difficulty ranking x and y directly treat the dominance relation between y and z as if it were a positive indicator of the value of y relative to x . Analogously in the case of the filling stations: if it is difficult to compare A and B, the unambiguous superiority of B to C provides a ‘reason’ for choosing B rather than A.

Another way of describing the asymmetric dominance effect is as a *shortlisting heuristic* for simplifying decision problems. One familiar way of limiting the cognitive demands of a decision problem is to begin by using ‘quick and dirty’ criteria to reduce the set of options to a more manageable shortlist. One such criterion, which has obvious utility if the aim is to reduce cognitive demands, is to eliminate options which would not be easy to compare with other items on the shortlist. (For example, consider an employer drawing up a shortlist of candidates for a job. Suppose that most candidates have at least adequate educational qualifications of some standard kind. If one candidate has non-standard and not obviously superior qualifications, that is likely to disadvantage her in the shortlisting process.) In effect, the asymmetric dominance heuristic selects a shortlist which, because its elements can be ranked by a dominance relation, makes decision-making particularly easy.

This shortlisting criterion is not as illogical as it might at first appear. Consider again the example of the filling stations. Suppose that all three prices have been drawn at random from the same unknown distribution (in which prices are expressed in, say, £/litre), and that a second, independent random process has determined whether each price is displayed as £/litre or as the £/gallon equivalent. If all you know in addition to this that B is cheaper than C, then B has the lowest expected price of the three. (The expected price at A is the mean of one random draw from the distribution; the expected price at B is the mean of the minimum of two random draws.)⁴

We suggest that, in markets in which tariffs can be spuriously complex, there is some tendency for consumers to use shortlisting heuristics which favour common standards. That is, when it is possible to do so, consumers are inclined to simplify their decision problems by eliminating offers that are expressed in individuated standards. In the example of the filling stations, such a heuristic would favour B (whose Imperial standard is shared with C) over A (whose metric standard is individuated).

Up to this point, we have not invoked any assumption that implies that prices expressed in common standards are systematically different from those expressed in individuated standards. We have argued that heuristics which favour common standards are psychologically salient, and that they can be ecologically valid even when the true value of an option is uncorrelated with the standard in which its price is expressed. We now show that there can be a Nash equilibrium for firms in which consumers' use of such a heuristic induces firms to charge competitive prices. Further, we show that, in *any* Nash equilibrium in which consumers' choices between tariffs satisfy a cognitively undemanding condition of 'price sensitivity', heuristics which favour common standards are utility-maximising.

3. The model

We present a model which shows the common standard effect at work in a very simple environment.⁵ For clarity in exposition, we adapt Jeffrey Perloff and Steven Salop's (1985) well-known model of a market with product differentiation. We focus on a market for a single consumer good, sold directly by producers to consumers. We consider the possibility that sellers might try to exploit the cognitive limitations of buyers by introducing spurious complexity into their pricing structures.

Although this case is chosen mainly for ease of modelling, it has practical interest in its own right. There are many examples, particularly in the telecommunications, electricity, gas and water industries, of markets in which firms compete to supply *exactly* the same product to consumers. In Britain, for example, domestic consumers can choose between competing electricity and gas suppliers, but the consumer has access to the same power and pipeline grids, irrespective of which supplier she chooses. In this environment, competition can only be in terms of prices. Suppliers typically offer a wide choice of tariffs, apparently catering to different patterns of electricity and gas use. There is evidence that consumers often fail to choose the lowest-cost supplier, which raises the possibility that tariff complexity is reducing competition (Wilson and Waddams Price, 2006).

Our formal model is of a one-period market for a good which is supplied by n competing firms, where $n \geq 3$.⁶ In interpreting the model, however, we imagine a sequence of periods in which the market is repeated, during which firms learn to follow optimal strategies and (more slowly) boundedly rational consumers gravitate towards

successful heuristics. There are N consumers, identical to one another ex ante. We assume that N/n is sufficiently large to legitimate the use of the law of large numbers when analysing the effects on firms of random variation at the level of the consumer. Each consumer buys a fixed quantity of the good, the same for all consumers; her problem is simply to satisfy this given demand at the lowest cost. As a normalisation, we define this quantity to be one *consumption unit* of the good. However, we do not assume that the consumer is consciously aware of this concept of quantity. As an example of the kind of situation to which our model might apply, consider a consumer who contracts with an electricity supplier to buy power according to a particular tariff over a fixed period, and then uses electricity as she needs it, without taking any account of the specificities of that tariff. At the end of the period, she is billed for whatever she has consumed (which, in fact, will be one ‘consumption unit’). The tariff might, for example, comprise a fixed charge, a charge per daytime kilowatt hour (kWh) and a charge per night-time kWh. In our terminology, a consumption unit is the consumer’s total consumption over the billing period, distributed between day and night according to the consumer’s pattern of electricity use (which, by assumption, is independent of the tariff). The consumer might have only a very hazy idea of how her consumption converts into kilowatt hours at different times of day, while being keenly aware, ex post, of the bottom line of the bill.

In this environment, there is scope for spurious complexity in tariff structures. Since purchases are the same for all consumers and are independent of the tariff under which they are bought, the relevant information in any tariff can be expressed as a single *price*, defined as the amount charged for one consumption unit. However, there are many different ways of presenting this information. As in our previous example, the unit might be subdivided into separately-priced components by using multi-part tariffs, or by charging different rates for consumption at different times of day. If two firms present their price information in sufficiently different forms, it may be difficult for consumers to work out which is offering the lower overall price. We represent this idea by modelling a tariff as a combination of a price and a *standard*. The price is an objective property of the tariff, about which the consumer is not directly informed. The standard is the device by which this information is presented. Any given standard is capable of expressing any given price. We will assume that each consumer’s ‘reading’ of any tariff is subject to random error; thus, each tariff provides only a noisy signal of its true price. However, if two tariffs use

the same standard, the consumer can make an accurate *ordinal* comparison of the corresponding prices.

As an illustration of the underlying idea, suppose that electricity tariff A has a fixed charge of £20 and a unit charge of £0.13 per kWh. Tariff B has a fixed charge of £10 per month and a unit charge of £0.14 per kWh. Tariff C has a fixed charge of £10 per month and a unit charge of £0.15 per kWh. Consider an individual whose average electricity consumption is 530 kWh/month. At this rate of consumption, a month's electricity will in fact cost £88.90 on tariff A, £84.20 on tariff B, and £89.50 on tariff C. If the individual is unsure about her rate of consumption, or is not good at arithmetic, her estimates of how much she would have to pay on each tariff may be subject to a good deal of error. However, it is easy for her to see that B is a cheaper tariff than C, even if she does not know *how much* less she will pay on B than on C. In the language of our model, the 'consumption unit' is 530 kWh/month, and the 'prices' charged by A, B and C are £88.90, £84.20 and £89.50. B and C are using a 'common standard' which allows ordinal comparisons of their prices to be made on a single dimension.⁷

In most real-world cases, different consumers, even when fully informed, may have different preferences over tariffs. In the case of multi-part electricity tariffs, for example, consumers whose demand is relatively low will prefer tariffs with low fixed charges and high rates per unit, while those with high demand will prefer the opposite. Thus, when consumers are differentiated, complexity in tariff structure can play a role in tailoring firms' offers to the tastes of individual consumers and in facilitating price discrimination. Even so, it remains true that complexity can make it harder for consumers to compare the offers of competing firms. The implication is that, from the viewpoint of consumers, there can be too much complexity and differentiation in tariff structures. Our modelling strategy allows us to isolate the component that is 'too much': the assumption that every consumer buys one consumption unit, irrespective of the tariff, allows the concept of 'spurious' complexity to be given a simple definition.⁸

In our model, each firm i has the same increasing and differentiable total cost function $C(q_i)$ where q_i is the firm's output, measured in consumption units. $C(\cdot)$ has a minimum efficient scale (MES) q^* , such that $q^* \leq N/n$ (so that, if consumer spending is distributed evenly between firms, all firms produce at or above MES). For $q_i \geq q^*$, $C(q_i) = cq_i$, where c represents both average and marginal cost. For $q_i \leq q^*$, average cost (AC) is

decreasing in quantity and marginal cost (MC) is non-decreasing; there are non-zero fixed costs, so that average cost tends to infinity as quantity tends to zero.⁹ There is no exit option for firms; thus each firm must incur at least the cost $C(0)$, even if its sales are zero. Figure 1 illustrates these assumptions.

In assuming a MES cost function, rather than one with constant average and marginal cost at all levels of output, we differ from Perloff and Salop and from many other models of markets with fixed numbers of firms. The latter type of cost function is often convenient in modelling because it implies a unique competitive price at which firms of different sizes can coexist. However, because it assumes the total absence of fixed costs, it has the unrealistic implication that a firm that sells a positive quantity, however small, at any price greater than marginal cost makes positive profit. Under this assumption, a competitive equilibrium would not be possible if *any* consumers could be induced by spurious complexity to pay more than the competitive price. It is more realistic to assume that if a firm unilaterally deviates from a putative competitive equilibrium by raising its price, it may earn *either* positive *or* negative profit, depending on the size of its market share after the price increase. An MES function has many of the convenient properties of a constant-cost function while allowing the effects of fixed costs to be modelled.

Firms seek to maximise expected profit. Each firm i sets a *tariff* (p_i, s_i) where p_i is its price per consumption unit and s_i is its standard; p_i is chosen from the set of strictly positive real numbers, and s_i from an infinite set S of possible standards. If the tariffs of any two firms i, j have the property that $s_i = s_j = s^*$, we will say that these firms use s^* as a *common standard*. If (and only if) i and j use a common standard, each consumer is informed of the true ranking of their prices. In consequence, each consumer has incomplete ranking information which can be represented by a partial ordering R of the set $\{p_1, \dots, p_n\}$. Notice that there can be more than one common standard in the market. A standard that is used by only one firm is *individuated*.

Each consumer's problem is to choose one (and only one) of these tariffs. For a representative consumer h , the ex post utility of choosing the tariff of firm i is $\alpha - p_i$, where α is the subjective value of a consumption unit supplied by any firm, normalised to monetary units. We implicitly assume that this value is sufficiently high that consumers always want to buy rather than not. If the ex post utility of each tariff was known to consumers ex ante, we would have a model of Bertrand competition. Instead, we assume

limitations on consumers' abilities to infer ex post utility from tariff information. We model these limitations as follows.

Consider any consumer h assessing the tariff of any firm i . We assume that, in addition to any information in the partial ordering R , the consumer receives a *price signal* r_{hi} where $r_{hi} = p_i + e_{hi}$. Here e_{hi} is an error term, representing the cognitive difficulty of inferring the price per consumption unit from the information provided by the tariff. We assume that e_{hi} is an iid random variable with zero mean, bounded support and a continuous and differentiable single-peaked density function.¹⁰

We assume a strategic interaction in which firms move first, simultaneously posting tariffs to which they are then committed. Next, each consumer chooses one of the posted tariffs; these choices determine the sales of each firm. Firms produce to meet these demands and incur the corresponding costs. Finally, consumers are billed for the quantities they have bought.

Given the tariffs chosen by firms, each consumer h receives the *information triple* $I_h = (s, r_h, R)$, where $s = (s_1, \dots, s_n)$ is the n -tuple of standards chosen by firms, $r_h = (r_{h1}, \dots, r_{hn})$ is the n -tuple of price signals received by h , and R is the incomplete ranking information. We define a *decision rule* for consumer h as a function f_h which, to each logically possible information triple I_h , assigns a probability distribution $f_h(I_h)$ over firms; this probability distribution determines the probability with which consumer h chooses the tariff of each firm, given the information I_h .

As the set of possible decision rules is very large, and as our topic is the implications for markets of consumers' bounded rationality, we will not assume that consumers use optimal decision rules. Our formal analysis will treat consumers' decision rules as exogenous. Given a specification of decision rules, our model describes a simultaneous-move game *for firms*, in which the set of pure strategies available to each firm i is the set of possible tariffs and in which payoffs are expected profits. We will investigate the Nash equilibria of such games. This analysis will be supplemented by a less formal discussion of the implications of assuming that, in a dynamic setting, consumers gradually learn to use decision rules that are expected-utility maximising, given the equilibrium behaviour of firms.

4. Decision rules

We focus on decision rules that are simple and psychologically salient. Consider the following two conditions that a decision rule f_h might satisfy:

Rational use of ranking information (RURI): If there are firms i, j such that $s_i = s_j$ and $p_i < p_j$, then the tariff of firm j is chosen with probability zero by consumer h .

Rational use of signal information (RUSI): If there are firms i, j such that s_i and s_j are individuated standards and $r_{hi} < r_{hj}$, then the tariff of firm j is chosen with probability zero by consumer h .

RURI is a simple condition with a self-evident rationale: if the consumer knows that firm i 's price is lower than firm j 's, there is no reason to buy from j . RUSI is clearly rational too, provided that the consumer has no information, additional to that in I_h , with which to distinguish between individuated standards. Given our interpretation of price signals as the consumer's best estimates of actual prices, RUSI seems psychologically salient: if the only perceptible difference between two tariffs is that one appears to have a lower price than the other, what reason can there be for choosing the one with the apparently higher price? Decision rules which satisfy RURI and RUSI will be called *price-sensitive*. We will restrict our attention to Nash equilibria in which consumers are price-sensitive.¹¹

The simplicity of the conditions RURI and RUSI reflects the fact that each of them operates on only one kind of price information (rankings and signals respectively). Optimally combining the two kinds of information is a difficult mathematical problem. At this point, one might expect consumers to rely on simplifying heuristics or algorithms. We now specify such an algorithm, based on a shortlisting operation of the kind discussed in Section 2:

Largest common standard (LCS): If there is at least one standard that is common to two or more firms, select the common standard that is used by the largest number of firms; then choose the lowest price associated with that standard. If all firms use individuated standards, choose the lowest price signal. Whenever there are ties (between standards, between prices or between signals), break them by randomisation.

Whenever possible, this algorithm reduces the decision problem to the largest (non-singleton) set from which a choice can be made by using ranking information. Because ranking information is used at the second stage, RURI is satisfied. Because the decision is

governed by price signals if the shortlisting heuristic cannot be applied, RUSI is satisfied. Thus, the LCS algorithm induces a price-sensitive decision rule.

We will show that this algorithm can support a Nash equilibrium in which all firms post the competitive price. However, it will be useful to compare LCS with two other simple and price-sensitive algorithms. The first of these is:

Dominance editing (DE): For every standard that is common to two or more firms, eliminate all firms using that standard except the one with the lowest price. Then choose the lowest price signal from among non-eliminated firms. Whenever there are ties, break them by randomisation.

This algorithm uses a psychologically salient *editing* operation – that of deleting dominated options before applying more sophisticated decision criteria – that has been identified experimentally in other decision contexts (Kahneman and Tversky, 1979).

While LCS favours common standards, DE is biased towards individuated standards. For example, take a case (such as the electricity tariffs of Section 3 or the filling stations of Section 2) in which there are three firms and two standards: one standard is unique to firm A, the other is common to firms B and C. Suppose that all three firms post the same price. LCS selects A, B and C with probabilities 0, 0.5 and 0.5 respectively. The corresponding probabilities for DE are 0.5, 0.25 and 0.25. The final algorithm we consider, although less well-grounded in experimental psychology than LCS and DE, is neutral between individuated and common standards:

Signal first (SF): Provisionally select the firm with the lowest price signal. Then restrict the choice to the set of firms which use the same standard as the provisionally selected firm; choose the lowest price associated with that standard. Whenever there are ties, break them by randomisation.

Applied to the example, SF would select each of A, B and C with probability 0.33.

5. IS equilibrium

In this section we investigate a class of Nash equilibria in which consumers are price-sensitive and all firms set individuated standards.

In the equilibria we consider, each firm sets its standard at random. This kind of randomisation seems to be essential for a credible equilibrium in which standards are

individuated. To see why, consider a putative equilibrium in which each firm sets a non-stochastic individuated standard. Assume for simplicity that all firms post the same price. Any firm can deviate from this putative equilibrium by selecting the same standard as another firm and marginally undercutting the latter's price. Given a credible specification of consumers' decision rules, such a deviation would increase the market share of the deviant firm. (For example, this would be the case if consumers used any of the algorithms LCS, DE or SF.) Thus, if the price in the putative equilibrium were above the competitive level c , any firm could increase its profit by a unilateral deviation. As we will show later, it is not compatible with equilibrium for individuated-standard firms to set prices at or below c .

We now define a class of *equilibria with individuated standards* or, for short, *IS equilibria*. In an IS equilibrium, each firm i chooses its standard s_i at random, independently of other firms. Since S is an infinite set, the probability that any two firms choose the same standard is zero. Each consumer h chooses to buy from the firm with the lowest price signal r_{hi} . Notice that this behaviour would be induced by any price-sensitive decision rule. The prices posted by the firms are such that no firm can increase its profit by unilaterally changing its price while continuing to randomise its standard. Given that all other firms randomise their standards, no firm can increase (or reduce) its profits by choosing its own standard deterministically, since the probability that a common standard will result remains zero. Thus, we have characterised a class of Nash equilibria in the game played by firms. In any such equilibrium, each consumer's behaviour is optimal, conditional on the information available to her.

A proof presented by Perloff and Salop (1985) can be adapted to show the existence of an IS equilibrium. Suppose that all firms post the same price p^I . Then (provided that the value of α is high enough to allow an internal solution) there is exactly one value of p^I that is consistent with IS equilibrium; this price is greater than c , which implies that firms make positive profits.¹² Here is an intuitive sketch of the proof.

The following is a necessary condition for any Nash equilibrium in which standards are individuated and firms' prices are non-stochastic: for any firm i , its price p_i must maximise its profit with respect to its *conjectural demand function* – that is, the function that plots how the quantity q_i sold by firm i varies with p_i when all other firms' prices remain unchanged. It is an elementary result in the theory of the firm that the marginal condition for profit-maximisation with respect to price is:

$$(1) \quad [p_i - C'(q_i)]/p_i = - (q_i/p_i)/(\partial q_i/\partial p_i).$$

The LHS of (1) is the *price-cost margin*; the RHS is the reciprocal of the *price elasticity of conjectural demand* (expressed as a positive number).

Suppose that all firms except one (say, firm j) set the price p^1 . Let $\eta(p^1)$ be the price elasticity of conjectural demand for firm j , expressed as a positive number and evaluated at $p_j = p^1$. Notice that when all firms set the same price, the quantity sold by each firm is N/n which, by assumption, is not less than MES; so, for each firm, marginal and average cost is c . Thus, adapting (1), p^1 is an equilibrium price if and only if

$$(2) \quad (p^1 - c)/p^1 = 1/\eta(p^1).$$

Now consider the determinants of $\eta(p^1)$. It follows from the specification of r_{hi} that, if all firms except j charge p^1 , the probability that j 's price signal is the lowest for any given consumer depends only (and negatively) on $p_j - p^1$. Thus, the *gradient* of j 's conjectural demand curve at $p_j = p^1$ is independent of p^1 . Since the quantity sold by j at $p_j = p^1$ is independent of p^1 (it is equal to N/n), the corresponding elasticity $\eta(p^1)$ is strictly positive and increasing in p^1 , tending to infinity as p^1 tends to infinity. Equivalently, the RHS of (2) is strictly positive and decreasing in p^1 , tending to zero as p^1 tends to infinity. Clearly, the LHS of (2) is increasing in p^1 , taking the value zero when $p^1 = c$ and tending to unity as p^1 tends to infinity. Thus (2) can be satisfied by one and only one value of p^1 ; this value is strictly greater than each firm's average cost c , implying positive profits.¹³

The foregoing argument establishes the existence of a unique IS equilibrium in which all firms post the same price. There may also be IS equilibria in which different firms post different prices, either deterministically or stochastically.¹⁴ However, it is easy to see that, in any equilibrium of this kind, all prices are greater than c . (A firm which posts the lowest of the equilibrium prices will sell at least the quantity N/n , implying a marginal cost of c . Because of (1), each firm's price is greater than its marginal cost, so the lowest price is greater than c .)

IS equilibrium can be interpreted as a state of affairs in which firms take advantage of consumers' cognitive limitations. Spurious complexity in tariffs prevents consumers from making accurate price comparisons. Because price signals are noisy, a firm can raise its price above the level charged by other firms while continuing to find buyers. This allows the market to support prices in excess of marginal and average cost.

6. CS equilibrium

We now investigate Nash equilibria in which consumers are price-sensitive and in which there are common standards. We begin by defining a particular equilibrium of this kind.

In an *equilibrium with a unique common standard* (or *CS equilibrium*), all firms use some common standard s^* and set the same price $p^C = c$. All consumers use the price-sensitive decision rule induced by the LCS algorithm. Thus, each firm sells the quantity N/n and makes zero profit. It is easy to see that this combination of strategies is a Nash equilibrium for firms. Any firm which deviates unilaterally from its equilibrium strategy earns strictly negative profit. If it posts a price greater than c while setting the standard s^* , it sells nothing because of the price-sensitivity of consumers, but incurs fixed costs. If it posts a price less than c while setting the standard s^* , it captures the entire market but sells at a loss. If it deviates from s^* , it sells nothing because consumers use the LCS algorithm, which shortlists only common-standard firms. Further, each consumer's behaviour is weakly optimal for her (trivially so, because there is no price variation).

The CS equilibrium can be interpreted as a state of affairs in which firms do *not* take advantage of consumers' cognitive limitations. Because firms use a single common standard, each consumer is able to make accurate ordinal comparisons between the prices posted by different firms. Thus, firms are in Bertrand competition; this induces marginal-cost pricing.

Clearly, marginal-cost pricing can be sustained in equilibrium only if a sufficiently large proportion of consumers use decision rules that favour common standards over individuated ones (as LCS does). If, to the contrary, most consumers used decision rules that were neutral between common and individuated standards (as SF is) or that favoured the latter (as DE does), a putative equilibrium in which all firms set a common standard and posted the competitive price would be vulnerable to unilateral deviation by a firm setting an individuated standard and posting a higher price. This is a corollary of the proposition (implied by the analysis in Section 5) that there is no IS equilibrium in which all firms post the competitive price.

Although CS equilibrium is not the only equilibrium in which consumers are price-sensitive and in which standards are not completely individuated, all such equilibria share the main characteristics of CS equilibrium. First, note that there may be price-sensitive

decision rules other than LCS which favour common standards over individuated ones and which can thereby support an equilibrium in which all firms use the same standard and post the competitive price. Another possibility is an equilibrium in which, although no firms set individuated standards, there are two or more common standards.¹⁵ In any such equilibrium, however, all firms post the competitive price. To see why, consider any putative equilibrium in which there are no individuated standards. For any given standard, all firms setting that standard are in Bertrand competition with one another, and so must post the same price, equal to the marginal cost of each of those firms. But, because firms are free to switch standards, equilibrium requires that all firms earn the same profit, irrespective of which standard they set.¹⁶ Because, by assumption, N/n (the average quantity produced by firms) is greater than or equal to q^* (MES), there must be at least one standard at which the marginal-cost price is c and profit is zero. Thus, *all* firms make zero profit. But, because of the properties of the cost function, any firm which makes zero profit while pricing at marginal cost must be posting a price equal to c .

7. Equilibrium with bargains and ripoffs

It remains to ask whether our model can have a Nash equilibrium in which consumers are price-sensitive and there is a mix of individuated and common standards. The answer is that such equilibria *are* possible; in all such equilibria, common-standard firms post the competitive price and individuated-standard firms post higher prices.

For simplicity, we show this only for the case in which all individuated-standard firms (the ‘IS sector’) post the same price p^I and sell the same quantity q^I , thus earning the same profit π^I , and in which all common-standard firms (the ‘CS sector’) set the same standard s^* , post the same price p^C and sell the same quantity q^C , thus earning the same profit π^C . (The argument can be generalised.)

If IS firms sell non-zero quantities in equilibrium, they face downward-sloping conjectural demand functions, as in IS equilibrium; thus p^I must be greater than IS firms’ marginal cost, $C'(q^I)$. If CS firms sell non-zero quantities in equilibrium, they are in Bertrand competition with one another; thus p^C must be equal to CS firms’ marginal cost, $C'(q^C)$. Since firms are free to move between sectors, equilibrium requires that $\pi^C = \pi^I$.¹⁷ Since firms in at least one sector must produce at or above N/n , there is at least one sector in which marginal cost equals c (compare the argument in the final paragraph of Section 6).

If this were true of the IS sector, IS firms would make positive profit (since then $p^I > C'(q^I) = C(q^I)/q^I$). But that is inconsistent with equilibrium, because profits in the CS sector cannot be positive. (In equilibrium, $p^C = C'(q^C)$; but, by the properties of the cost function, $C'(q^C) \leq C(q^C)/q^C$.) So CS firms produce at or above N/n . Then these firms have marginal cost equal to c , implying $p^C = c$ and $\pi^C = 0$. Equilibrium then requires $\pi^I = 0$. Since IS firms price above marginal cost, this is possible only with $q^I < q^*$ and $p^I > c$.

In an equilibrium of this kind, as in Salop and Stiglitz's (1977) model of monopolistically competitive price dispersion, there are 'bargains' (CS tariffs) and 'ripoffs' (IS tariffs). Clearly, such an equilibrium could not be sustained if consumers were fully rational, because a firm's use of a common standard is an observable signal indicating that its price is at the competitive level. Nor could it be sustained in a market in which all consumers used the LCS decision rule, since that rule would shortlist only CS firms, leaving IS firms with zero sales and negative profit. But it *could* be compatible with some mixtures of price-sensitive decision rules. For example, suppose that some consumers use the LCS rule while others use DE and/or SF. LCS consumers buy only from CS firms, but some DE and SF consumers pay the higher prices of IS firms. Depending on the relative frequencies of the different decision rules, there could be an equilibrium distribution of firms between the two sectors such that $\pi^I = \pi^C = 0$.¹⁸

8. Learning by consumers

We have shown that three types of Nash equilibrium (and only these) are consistent with the assumption that consumers are price-sensitive. In one type of equilibrium, exemplified by IS equilibrium, spuriously complex tariffs allow prices to be maintained above the competitive level. In a second type, exemplified by CS equilibrium, common standards induce Bertrand competition and hence competitive prices; such an equilibrium can be maintained only if a sufficiently large proportion of consumers use decision rules which favour common standards. In the third type of equilibrium, common-standard firms posting competitive prices coexist with individuated-standard firms posting higher prices; such a 'bargains and ripoffs' equilibrium is possible only if some consumers favour common standards and some do not.

Notice that in all these equilibria, firms which use common standards post competitive prices, while firms which use individuated standards post prices that are higher

than the competitive level. Thus, in *every* equilibrium it is expected-utility-maximising for consumers to buy only from common-standard firms if any such firms exist; in *some* equilibria, this behaviour is strictly optimal. The implication is that heuristics which favour common standards are not only psychologically salient; they are also well-adapted to environments in which a variety of equilibria can be encountered. Markets for consumer goods, considered in general, might be thought of as such an environment.

The idea of favouring common standards is very general. It is not tied to any specific standard, to any specific firm or firms, or to any specific type of product. Thus, it might be learned in one context and then applied in others.¹⁹ If consumers are boundedly rational, gradually learning to follow utility-maximising heuristics, and if there is a general tendency for common standards to be associated with low prices, heuristics which favour common standards will tend to emerge.

A complementary possibility is that *disequilibrium* behaviour by firms may support the process by which consumers learn to favour common standards. The concept of completely randomised standards, as used in the definition of IS equilibrium, is a modelling simplification. The nearest realistic equivalent to randomisation is a situation in which each firm changes its standard frequently and unpredictably and, when doing so, avoids standards that are currently used by other firms. If changing standards is costly, or if there is some constraint on the frequency with which changes are made, the choice of standards becomes a game of strategy between firms. A crucial component of such a game is the fact that, if two or more firms are pricing above the competitive level and if one firm (say i) can predict the price and standard that another firm (j) will set in a given period, then i can gain sales at j 's expense by replicating j 's standard while undercutting its price. It is this possibility of being undercut that forces firms to keep changing standards, and the constant change in standards is essential for the sustainability of non-competitive prices. But if firms are to continue to respond to the possibility of being undercut, the probability of being undercut must be non-zero. Thus, individuated standards and non-competitive prices can persist only in combination with episodes of undercutting. In other words: in a realistic form of IS equilibrium, individuated standards will be the norm, but there will be occasional episodes of price competition between firms which are temporarily using common standards. A market with these characteristics will provide consumers with evidence of the association between common standards and low prices.

This analysis also suggests how a transition from IS to CS equilibrium might come about. A realistic form of IS equilibrium will be a state of turbulence in which firms are constantly changing standards for tactical purposes, sometimes with the intention of finding a standard that is unique to themselves, but sometimes with the contrary intention of replicating other firms' standards and competing on price. The more frequent are such episodes of price competition, the greater is the incentive for consumers to use decision rules which favour common standards. But the greater is the proportion of consumers who use such rules, the greater is the incentive for firms to undercut one another. (A firm that undercuts a rival not only takes sales from the rival, but also attracts consumers who favour common standards.) Thus, starting from an equilibrium with IS characteristics, tactical manoeuvres by firms might initiate a chain reaction leading to CS equilibrium. One implication of this argument is that the long-run survival of individuated standards may require collusion among firms, with some form of tacit agreement that firms do not replicate one another's standards. (Conversely, a transition from CS to IS equilibrium might be brought about by collusion, if sufficiently many firms individuated their standards in concert.)

9. Conclusion

Our tentative conclusion is that there may be general market mechanisms which, over the long run, favour the evolution of common standards. Spuriously complex tariffs may be either a transitory phenomenon in the evolution of markets for particular goods, or symptomatic of tacit collusion.

We present this conclusion as a contribution to the understanding of markets, and not as an argument against regulation. To the contrary, our analysis can be read as a rationale for some degree of light-touch regulation to impose common standards on tariff structures. Such regulation is 'light-touch' in the sense that it supports a transition from one Nash equilibrium (with high prices) to another (with low prices); once the transition is complete, and provided that firms are not able to collude, the regulation may be self-enforcing. Expressing the qualification about collusion in a different way, the imposition of common standards might be interpreted as an anti-collusion policy. Further, it may not be necessary that all firms are regulated; all that is needed is that the number of firms that are required to use a common standard is enough to initiate a process of transition and to

ensure that, once an common-standard equilibrium has become established, unregulated firms cannot increase their profits by collusively switching to individuated standards.

Nevertheless, we believe that any discussion of regulation should take account of the self-regulating powers of the market system. One should be cautious about inferring, from the growing evidence of the cognitive limitations of economic agents, that unregulated markets necessarily overburden consumers' decision-making capacities. It is important to consider how decision heuristics and pricing conventions might evolve to help boundedly rational consumers navigate the complexity of the market.

References

- Camerer, Colin, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue and Matthew Rabin (2003). Regulation for conservatives: behavioral economics and the case for 'asymmetric paternalism'. *University of Pennsylvania Law Review*, 151: 1211-1254.
- DellaVigna, Stefano and Ulrike Malmendier (2004). Contract design and self-control: theory and evidence. *Quarterly Journal of Economics* 119: 353–402.
- DellaVigna, Stefano and Ulrike Malmendier (2006). Paying not to go to the gym. *American Economic Review* 96: 694–719.
- Dudley, Marc (1990). Competition by choice: the effect of consumer search on firm location decisions. *American Economic Review* 80: 1092–1104.
- Eliaz, Kfir and Ran Spiegler (2006). Contracting with diversely naive agents. *Review of Economic Studies* 73: 689–714.
- Ellison, Glenn (2005). A model of add-on pricing. *Quarterly Journal of Economics* 120: 585-637.
- Ellison, Glenn and Sara Fisher Ellison (2009). Search, obfuscation, and price elasticities on the Internet. *Econometrica* 77: 427–452.
- Gabaix, Xavier and David Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121: 505-540.
- Gaudeul, Alexia and Robert Sugden (2007). Spurious complexity and common standards in markets for consumer goods. CCP Working Paper No. 07-20, Centre for Competition Policy, University of East Anglia.
- Grubb, Michael (2008). Selling to overconfident consumers. Working paper, forthcoming in *American Economic Review*.
- Heidhues, Paul and Botond Köszegi (2004). The impact of consumer loss aversion on pricing. Working paper SP II 2004–17, Social Science Research Centre, Berlin.
- Huber, J., J.W. Payne and C. Puto (1982). Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *Journal of Consumer Research* 9, 90–98.

- Kahneman, Daniel and Amos Tversky (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47: 263-291.
- Marks, Peter (2002). *Association between games: a theoretical and empirical investigation of a new explanatory model in game theory*. Amsterdam: ThelaThesis
- Perloff, Jeffrey and Steven Salop (1985). Equilibrium with product differentiation. *Review of Economic Studies* 52: 107-120.
- Shafir, Eldar, Itamar Simonson and Amos Tversky (1993). Reason-based choice. *Cognition* 49: 11-36.
- Salop, Steven and Joseph Stiglitz (1977). Bargains and ripoffs: a model of monopolistically competitive price dispersion. *Review of Economic Studies* 44: 493–510.
- Shafir, Eldar, Itamar Simonson and Amos Tversky (1993). Reason-based choice. *Cognition* 49: 11-36.
- Shafir, S., T.A. Waite and B.H. Smith, B. H. (2002). Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioral Ecology and Sociobiology*, 180–187.
- Spiegler, Ran (2006). Competition over agents with boundedly rational expectations. *Theoretical Economics* 1: 207–231.
- Sugden, Robert (2004a). The opportunity criterion: consumer sovereignty without the assumption of coherent preferences. *American Economic Review* 94: 1014-1033.
- Sugden, Robert (2004b). *The Economics of Rights, Co-operation and Welfare*, second edition. Palgrave Macmillan.
- Sugden, Robert (2009). Market simulation and the provision of public goods: a non-paternalistic response to anomalies in environmental evaluation. *Journal of Environmental Economics and Management* 57: 87–103.
- Sunstein, Cass R. and Richard H. Thaler (2003a). Libertarian paternalism is not an oxymoron. *University of Chicago Law Review*, 70: 1159-1202.
- Sunstein, Cass R. and Richard H. Thaler (2003b). Libertarian paternalism. *American Economic Review, Papers and Proceedings*, 175-179.

Thaler, Richard H. and Cass R. Sunstein (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.

Wilson, Chris and Catherine Waddams Price (2006). Do consumers switch to the best supplier? Discussion Paper, University of East Anglia and ESRC Centre for Competition Policy.

Figure 1: Implications of the MES assumption

Notes

¹ At the time of writing, there are some indications that long-standing conventions about add-ons in the airline industry are breaking down. Very recently, ‘full-service’ airlines have started to introduce new add-on charges, for example for checked baggage. It is too early to know whether these new practices will become the norm. However, the fact that these add-ons are being introduced by different airlines at almost the same time is suggestive of tipping from one equilibrium to another (or of tacit collusion). This is consistent with the hypothesis that the previous common standard was a convention, even if not a fully stable one.

² http://www.travelodge.co.uk/common_room_questions

³ In fact, 1 Imperial gallon = 4.5454 litre, so B is cheapest.

⁴ Our problem bears some resemblance to the Monty Hall problem. Imagine a game show in which the contestant chooses one of three closed boxes, and wins the contents of that box. Suppose that the value of each box has been drawn independently from a binary distribution over the values £0 (with probability 2/3) and £1000 (with probability 1/3). The show's host, who knows the value of each box, selects two boxes at random and then points to one of these, informing the contestant that it is at least as valuable as the other. (If they are equally valuable, the one he points to is selected at random.) Then, by the reasoning in the main text, the contestant should open the pointed-to box. Now suppose this game is changed so that exactly one box contains £1000 (i.e. the random draws are not independent) and the contestant (rather than a random mechanism) selects the two boxes for which the host will provide ranking information. It is still optimal to choose the pointed-to box. But the revised problem is just a re-framing of the Monty Hall problem. In the Monty Hall problem, the contestant selects the two boxes for which information will be given by ‘provisionally choosing’ the remaining box; the host provides the ranking information by opening an empty box; the contestant is then invited to ‘switch’ boxes.

⁵ There are some similarities between our model of pricing standards and Dudley’s (1990) model of the location choices of retail firms. In our model, common pricing standards both induce and signal Bertrand competition. In Dudley’s model, common spatial locations both induce and signal Cournot competition.

⁶ Some of our results would not hold for a market with only two firms. In our model, the demand conditions for a firm which shares a standard with at least one other firm are different from those for a firm whose standard is unique to itself. Our analysis of ‘CS equilibrium’, in which all firms use the same standard, relies on the property that if one firm deviated from that standard, the other firms would still share a standard.

⁷ For modelling purposes, we assume that each tariff has one and only one standard. This implies that the set of tariffs can be partitioned into subsets, each of which is defined by a single standard; accurate ordinal comparisons can be made between tariffs if and only if they belong to the same subset. We neglect the complication that, in reality, the sets within which ordinal comparisons can be made might overlap.

⁸ In the simple case represented by our model, *all* complexity is spurious. Thus, a regulation which imposed a common standard could not harm consumers (although it might reduce firms’ profits). Even if the case for such a regulation seems uncontroversial, this should not be seen as an objection to our model. Our aim is to investigate whether common standards can emerge and be sustained without regulation.

⁹ Cost functions of the MES type are generated if there are constant returns to scale in production but the firm has fixed costs in the form of a commitment to buy a minimum vector of inputs. As a simple example, let q be output and let l_k be the quantity of input k used in production ($k = 1, \dots, m$). Assume a Leontief production function $q = \min_k(l_k/a_k)$, where $(a_1, \dots, a_m)$ is a vector of positive coefficients. For each input k there is a unit price w_k and a minimum quantity l'_k which the firm is committed to buying. Thus, expenditure on each input k is $w_k \max(l'_k, l_k)$. This gives a piecewise linear MES total cost function.

¹⁰ By assuming that the distribution of e_{hi} is the same for all tariffs, we abstract from the possibility that some standards are more difficult to understand than others. Our hunch is that people’s intuitive sense of ‘simplicity’ in tariffs, product specifications, labelling, and so on is often a matter of convention: it is easier to process information if it comes in familiar forms. Our concern in this paper is with the emergence of common standards, not with ‘intrinsic’ simplicity.

¹¹ Thus, we do not discuss the trivial equilibrium in which all firms post the monopoly price α and consumers choose between firms at random, ignoring price signals. In this equilibrium, consumers’ behaviour is weakly optimal but not price-sensitive.

¹² Perloff and Salop assume that each consumer h chooses the firm i such that $\theta_{hi} - p_i$ is maximised, where p_i is the price charged by firm i and θ_{hi} is a random variable representing the ‘value’ of firm i ’s product to consumer h . The latter variable plays the same role as $\alpha - e_{hi}$ in our model.

¹³ We are assuming that (2) is satisfied at a value of p^I less than α , permitting an internal solution. If this is not the case, there is a corner solution at the monopoly price $p^I = \alpha$.

¹⁴ We are unable to prove any general result about the existence or non-existence of such equilibria. Perloff and Salop (1985) report that they are unable to prove corresponding general results for their model.

¹⁵ This would not be possible if all consumers used the LCS rule, since that rule favours the common standard that is used by the largest number of firms. However, an equilibrium could be supported by other decision rules that favour common standards over individuated ones.

¹⁶ Here we abstract from discontinuities caused by the lumpiness of firms’ market shares. Strictly, if the number of firms using each standard is small, the Nash equilibrium condition that no firm can increase its profit by changing standards implies only that cross-standard profit differences are small.

¹⁷ Again, we abstract from problems of discontinuity. Compare note 14.

¹⁸ For further analysis of this case, see Gaudeul and Sugden (2007).

¹⁹ Compare Sugden’s (2004b, pp. 49-54) discussion of how conventions can spread from one context to another by analogy: rules which have more general application and are more susceptible to analogy are better equipped to reproduce themselves. See also Marks (2002).