



Spurious Complexity and Common Standards in Markets for Consumer Goods*

by

Alexia Gaudeul

Dept. of Economics and Centre for Competition Policy, University of East Anglia

&

Robert Sugden

Dept. of Economics, University of East Anglia

CCP Working Paper 07-20

Abstract: Behavioural and industrial economists have argued that, because of cognitive limitations, consumers are liable to make sub-optimal choices in complex decision problems. Firms can exploit these limitations by introducing spurious complexity into tariff structures, weakening price competition. This paper models a countervailing force. Consumers' choice problems are simplified if competing firms follow common conventions about tariff structures. Because such a 'common standard' promotes price competition, a firm's use of it signals that its products offer value for money. If consumers recognize this effect, there can be a stable equilibrium in which firms use common standards and set competitive prices.

November 2007

JEL Classification Codes: D83; L13; L15; L51

Keywords: decision-making; naïve consumers; savvy consumers; price competition; common standard effect; cognitive limitations.

Acknowledgements:

* 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, and at seminars at the University of East Anglia and the University of Sheffield.

We thank participants in these meetings for their comments and suggestion. The support of the Economic and Social Research Council is also gratefully acknowledged.

Contact details:

Alexia Gaudeul, Dept. of Economics and Centre for Competition Policy,
University of East Anglia, Norwich, NR4 7TJ, UK.

a.gaudeul@uea.ac.uk

Robert Sugden, Dept. 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 paternalistic interventions to simplify consumers' choice problems, by imposing what are claimed to be only mild restrictions on their freedom of choice (Sunstein and Thaler, 2003a, 2003b; Camerer et al, 2003). 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 make markets less competitive, inducing welfare losses (Ellison and Ellison, 2004; Ellison, 2005b; Gabaix and Laibson, 2006; Spiegler, 2006). 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 which is intrinsic to competitive markets: 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. By facilitating comparisons between products, such conventions 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. Thus, consumers can learn by experience to favour products which meet common standards. If consumers act in this way, profit-seeking firms are induced to adopt those standards and are penalised for deviating from them.

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*.

In this paper, 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 evolve and become self-sustaining.

1. The intuition: does Wal-Mart offer too much choice?

As an introduction to the intuition behind our argument, consider the following passage from the paper in which Cass Sunstein and Richard Thaler (2003a) advocate 'libertarian paternalism':

How much choice should people be given? Libertarian paternalists want to promote freedom of choice, but they need not seek to provide bad options, and among the set of reasonable ones, they need not argue that more is necessarily better. Indeed that argument is quite implausible in many contexts. In the context of savings plans, would hundreds of thousands of options be helpful? Millions? Thirty years ago, most academics had only two investment options in their retirement plan, TIAA and CREF. Now most universities offer more than one provider and often dozens, if not hundreds, of funds from which to choose. ... Do participants gain from this increase in their choice set? ... [O]ne recent study finds that when [US pension] plans offer more choice, participants are slower to join, perhaps because they are overwhelmed by the number of choices and procrastinate. (pp. 1196-1197)

The suggestion is that, even when there is no spurious complexity, unregulated markets can present consumers with too much choice, and that there can be a case – indeed, a case that could be accepted by libertarians – for paternalistic interventions designed to reduce the range of choice.

Sunstein and Thaler’s paper has sparked off a vigorous debate on the web. One participant quotes Newt Gingrich as having said:

If you were to walk into a Wal-Mart and say to people, ‘Don’t you feel really depressed by having 258,000 options; shouldn’t it be their obligation to reduce the choice you must endure?’, they would think you were nuts.¹

Gingrich is surely right: most supermarket customers *would* be astonished at the suggestion that the range of choice presented to them was too large. Sunstein and Thaler implicitly recognise the implausibility of the ‘too much choice’ claim in relation to ordinary consumer purchases by restricting its application to situations in which consumers are poorly informed or lacking in experience; they allow that ‘better informed choosers can more easily navigate the menu options’ (pp. 1197-1198). Thus, they might accommodate Gingrich’s objection by arguing that the typical consumer has well-informed preferences about the goods in Wal-Mart, while the typical employee lacks such preferences about pension plans. But is this kind of appeal to informed preferences sufficient to account for consumers’ confidence in navigating supermarkets?

Imagine a store which stocks the 258,000 Wal-Mart options, but in which these goods are arranged on the shelves in a random order, changed every 24 hours. Further, imagine that all products are packaged in plain white containers; on each package, the nature of its contents is described in black print in a standard typeface. We conjecture that if consumers had no choice but to shop at such a store, they would find shopping an extremely onerous task, and would welcome a reduction in the number of options. The point of this thought experiment is that our ability to navigate supermarkets is highly

¹ The quote is from a discussion piece entitled ‘Choice and its enemies’, dated 19 April 2005, by Pejman Yousefzadeh, at <http://www.tcsdaily.com/article.aspx?id=041905D>.

dependent on the existence of conventions about how options are displayed. One such set of conventions governs retailers' decisions about which products are placed close to which. For example, in just about all supermarkets, the different coffee products are placed close together, and relatively close to the different tea products. A customer who is looking for tea knows she is in roughly the right part of the store when she sees coffees; when she locates the tea section, she can readily compare the different teas. Another set of conventions governs producers' decisions about the packaging of their products. For example, there are family resemblances among the package designs used by different tea producers. Because of these features, the customer can quickly locate tea products against a background of other groceries. It seems undeniable that there must be some mechanisms at work in retail markets, favouring the emergence and persistence of conventions that reduce the complexity of consumers' choice problems. This paper investigates one such mechanism.

Of course, casual observation also reveals many cases in which producers and retailers contravene established conventions as a profit-seeking strategy. For example, supermarkets sometimes place their 'special offers' away from the shelves used to display similar but normally-priced goods. It would be naïve to deny that in many such cases, retailers are seeking to exploit consumers' cognitive limitations. However, competition surely restricts the scope for this kind of obfuscation. If consumers find it easier to get value for money when they shop in supermarkets which use standardised layouts, they will tend to patronise supermarkets that are laid out in standard ways; retailers who try to entrap customers by using unfamiliar layouts will lose business. Intuitively, it seems that we are observing a balance of forces, some of which favour the emergence of common standards while others favour 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 model

We present a model which shows the common standard effect at work in a very simple environment. 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).

The case we model has more general theoretical interest. Robert Sugden (2004a) has investigated how far competitive markets can deliver normatively desirable outcomes when consumers lack well-defined and consistent preferences. He presents a model in which, for the market to be efficient in generating opportunities for consumers, it is sufficient that there is free entry for profit-seeking arbitrageurs, and that consumers are 'price-sensitive'. At any given moment, a price-sensitive consumer buys goods only at the lowest prices currently quoted for them (and sells only at the highest); in all other respects, she may be highly irrational (for example, by buying a good at a high price and then immediately selling it at a lower price). Clearly, this result is blunted if, in reality, the task of finding the most favourable price is cognitively demanding. Sugden's model implicitly assumes that all prices are quoted in a standard form, so that price comparisons can be made easily. The present paper investigates whether such a standard might be self-enforcing.

Our model is of a one-period market for a good which is supplied by n competing firms, where $n \geq 3$.² Formally, we investigate equilibrium properties of this market. In interpreting the model, however, we imagine a sequence of periods in which the market is repeated, during which firms and consumers learn to follow optimal strategies. Implicitly, we assume that consumers cannot recognise the continuing identities of individual firms through time; this screens out of the model any effects of reputation-building. Our equilibrium is to be interpreted as the end-point of a learning process.

There are N consumers, identical to one another *ex ante*. (The model includes some stochastic variables whose realisations may differ between consumers.) 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.

² Many 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.

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 £10 per month, a day rate of £0.14 per kWh, and a night rate of £0.06 per kWh. Tariff B has the same fixed charge, the same definitions of 'day' and 'night', and the same night rate, but a day rate of £0.12 per kWh. It is easy for a consumer to see that tariff B is cheaper, even if she does not know how much it will cost her to buy electricity on each tariff, or how much less she will pay on tariff B than on tariff A. In the language of our model, the two tariffs are using a 'common standard' (one in which differences in price are expressed through differences in the day rate).

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. 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'.

In our model, each firm i has the same 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 is decreasing in quantity and marginal cost is non-decreasing; there are non-zero fixed costs, so that average cost tends to infinity as quantity tends to zero.³

Each firm seeks 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 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*. Notice that there can be more than one common standard in the market. A standard that is used by only one firm is *individuated*.

³ We differ from Perloff and Salop, and from many other models of markets with fixed numbers of firms, by assuming a MES cost function rather than one with constant marginal cost at all levels of output. The latter type of cost function has the unrealistic property that a firm which sells a positive quantity at any price greater than marginal cost can make positive profit, no matter how small its sales. We will be analysing cases in which some firms are in Bertrand competition, pricing at marginal cost, while other firms have the option of setting higher prices while still having positive sales. We want to leave open the question of whether (or under what circumstances) this option is profitable. 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.

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 $\alpha + \beta v_{hi} - p_i$, where $\alpha + \beta v_{hi}$ is the subjective value of a consumption unit supplied by firm i , normalised to monetary units. The parameter $\alpha > 0$ represents the average value of a consumption unit, ranging across all consumers and all firms; we implicitly assume that this value is sufficiently high that consumers always want to buy rather than not. We use βv_{hi} to represent *idiosyncratic value* – that is, a component of subjective value that is specific to the match between a particular consumer and a particular firm. The term v_{hi} is an iid random variable with zero mean, bounded support and a continuous and differentiable single-peaked density function. The parameter $\beta > 0$ allows us to formalise the idea that idiosyncratic value is a very small component of subjective value; where appropriate, we do this by taking limits as $\beta \rightarrow 0$.⁴

If the *ex post* utility of each tariff was known to consumers *ex ante*, we would have a model very similar to that of Perloff and Salop. Crucially, however, 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 the consumer receives a signal r_{hi} , where

$$(1) \quad r_{hi} = \alpha + \beta v_{hi} - p_i + e_{hi}.$$

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.⁵

⁴ For modelling purposes, it is convenient to assume a small component of idiosyncratic value. As will emerge later, this assumption ensures that the demand elasticity for common-standard firms is finite. It also ensures that it is optimal for consumers to use signals when comparing tariffs with different standards, even in equilibria in which all tariffs have the same price.

⁵ 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. Whatever the truth of this, our concern in this paper is with the emergence of common standards, not of 'intrinsic' simplicity.

If every firm uses a different standard, the n signals constitute the whole information available to the consumer. Clearly, her optimal strategy is then to choose the tariff with the highest signal. (Because of the continuity assumptions we have made about idiosyncratic value and error, ties occur with zero probability and so can be ignored.) If, however, some firms use common standards, the consumer receives additional information. For any firms i, j , if $s_i = s_j$, then the consumer is informed of the true ranking of the corresponding *ex post* utilities – that is, she knows the sign of $[\beta v_{hi} - p_i] - [\beta v_{hj} - p_j]$.⁶ Thus, if all firms use the same standard, the optimal strategy is to choose the tariff with the highest true ranking, independent of the signals. Since some significant properties of our model can be shown by using only these two obvious optimality principles, we postpone the question of how the consumer should choose when at least some firms use common standards, but it is not the case that all firms use the same standard.

3. Two equilibria

Our concept of equilibrium is Bertrand–Nash. That is, we assume a strategic interaction in which firms move first, simultaneously choosing tariffs; next, each consumer chooses a tariff, determining the sales of each firm; firms produce to meet these demands and incur the corresponding costs; and finally, consumers are billed for the quantities they have bought.⁷ A state of the model can be described by specifying the tariff (p_i, s_i) of each firm $i = 1, \dots, n$. For this state to be an equilibrium, it must be the case that no firm can increase its profit by changing its tariff, given the tariffs of other firms, and given the decision rules used by consumers. It must also be the case that, given the tariffs set by the firms, each consumer's decision rule maximises her expected utility.

⁶ This specification implies that, within a given standard, the consumer can integrate differences of idiosyncratic value into her ranking of tariffs. In terms of our illustrative story of electricity tariffs, it might be thought more realistic to assume that the within-standard information is the *sign* of $p_i - p_j$, and that this has to be evaluated in combination with the *numerical value* of $\beta(v_{hi} - v_{hj})$, i.e. the difference in idiosyncratic value. Given that idiosyncratic value plays only a minor role in our model, we prefer to use the simpler specification of the main text.

⁷ The assumption of an MES cost function implies that there are no capacity constraints and that the profit of firm i is strictly increasing in q_i at any given price $p_i > c$, and non-decreasing at $p_i = c$. In the equilibria we consider, $p_i \geq c$ for every firm i . Thus, no firm would want to produce less than the quantity it can sell at its posted price.

Thus, the following is a necessary condition for equilibrium: for any firm i , holding constant its standard s_i , its price p_i must maximise its profit with respect to its (Bertrand–Nash) *conjectural demand function* – that is, the function that plots how the quantity q_i sold by firm i varies with p_i and s_i when all other firms' tariffs 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:

$$(2) \quad [p_i - C'(q_i)]/p_i = - (q_i/p_i)/(\partial q_i/\partial p_i).$$

The LHS of (2) is the *price-cost margin*; the RHS is the reciprocal of the *price elasticity of conjectural demand* (expressed as a positive number). This result will be used repeatedly in our analysis.

We now characterise two equilibria. The first is *equilibrium with individuated standards* or, for short, *IS equilibrium*. In IS equilibrium, each firm i chooses its standard s_i at random. Since S is an infinite set, the probability that any two firms choose the same standard is zero. All firms set the same price, p^l . Each consumer h chooses to buy from the firm with the highest signal r_{hi} . Adapting a proof presented by Perloff and Salop, there is exactly one value of p^l that is consistent with Bertrand–Nash equilibrium; this price is greater than c , which implies that firms make positive profits.⁸

Here is an intuitive sketch of the proof. First notice that, given that all firms use different standards, consumers optimise by choosing between tariffs according to their signals. Given that consumers choose in this way, firms' sales and hence their profits are independent of the standards they use. So, in proving the existence of IS equilibrium and in investigating its comparative-static properties, we need make no further reference to standards. Suppose that all firms except one (say, firm j) set the price p^l . Let $\eta(p^l)$ be the price elasticity of conjectural demand for firm j , expressed as a positive number and evaluated at $p_j = p^l$. Notice that when all firms set the same price, the quantity sold by each firm is N/n which, by assumption, is more than MES; so, for each

⁸ 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 $\alpha + \beta v_{hi} + e_{hi}$ in our model.

firm, marginal cost is c . Thus, adapting (2), p^l is an equilibrium price if and only if

$$(3) \quad (p^l - c)/c = 1/\eta(p^l).$$

Now consider the determinants of $\eta(p^l)$. It follows from the specification of r_{hi} in (1) that, if all firms except one (say, firm j) charge p^l , the probability that j 's signal is the highest for any given consumer depends only (and negatively) on $p_j - p^l$. Thus, the *gradient* of j 's conjectural demand curve at $p_j = p^l$ is independent of p^l . Since j 's sales at $p_j = p^l$ are independent of p^l (they are equal to N/n), the corresponding elasticity $\eta(p^l)$ is strictly positive and increasing in p^l . Equivalently, the RHS of (3) is strictly positive and decreasing in p^l . Clearly, the LHS of (3) is increasing in p^l , taking the value zero when $p^l = c$ and tending to infinity as p^l tends to infinity. Thus (3) can be satisfied by one and only one value of p^l .

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 cost, even in the absence of 'genuine' product differentiation (that is, even when $\beta \rightarrow 0$).

The second equilibrium is *equilibrium with a common standard*, or *CS equilibrium*. Consider the decision rule followed by any given consumer in choosing between tariffs. We will say that a decision rule *rejects individuated standards* if it chooses a tariff with a common standard whenever at least one such tariff is available. A decision rule *respects ordinal information* if, whenever there is information about the true ranking of the *ex post* utilities of two tariffs (that is, whenever those tariffs have a common standard), it never chooses a tariff which is known to be inferior. In CS equilibrium, all consumers follow decision rules which respect ordinal information; for a sufficiently high proportion of consumers, their decision rules also reject individuated standards. All firms use the common standard s^* and set the

same price p^C . Given that all firms behave in this way, the decision rules used by consumers are weakly optimal for them. (It is both necessary and sufficient for optimality that decision rules respect ordinal information.) Adapting the same proof as before from Perloff and Salop, there is exactly one value of p^C that is consistent with Bertrand–Nash equilibrium, given that all firms use a common standard. As $\beta \rightarrow 0$, the equilibrium value of p^C tends to c and profits tend to zero. Given that a sufficiently high proportion of consumers use rules which reject individuated standards, no firm can benefit by deviating unilaterally from the common standard. (If *all* consumers use rules of this kind, any firm which deviates from the common standard sells nothing and so makes negative profit.)

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 *ex post* utility of buying from different firms. As the degree of product differentiation (represented by β) tends to zero, the relationship between firms converges to Bertrand competition, and price tends to marginal cost.

A comparison between IS and CS equilibrium suggests that the strategy of rejecting non-shared standards is self-validating at the market level. Intuitively, that strategy can be rationalised in terms of a belief that shared standards are a signal of low prices. If all consumers follow this strategy, an equilibrium can be sustained in which firms use a market-wide common standard; and the existence of such a common standard induces low prices. This gives some reason to expect that CS equilibrium is stable. We will firm up that intuition by investigating the dynamics of our model.

4. Dynamics

For the purposes of our dynamic analysis, we treat the model as described in Section 2 as a game that is played repeatedly by the n firms and N consumers. Over time, firms revise their tariffs (by adjusting prices,

standards, or both) in the direction of increased profitability, while consumers revise their decision rules in the direction of increased expected utility.

The analysis in this section will be somewhat speculative. It is extremely difficult to prove even static equilibrium results, even for the apparently simple Perloff–Salop model in which marginal costs are constant and the issue of common standards is not considered. For example, it seems natural to expect that, in the Perloff–Salop model, an increase in the number of firms would induce a fall in the equilibrium price; but this cannot be proved in general. And it is not known whether that model has equilibria in which different firms charge different prices (Perloff and Salop, 1985). We are investigating the dynamic properties of a considerably more complicated model.

Our dynamic analysis is based on the following simplifying assumptions:

Minimal idiosyncratic differentiation. As in the analysis of IS and CS equilibria, we consider only the limiting case of the model as $\beta \rightarrow 0$.

Naïve and savvy consumers. Since the purpose of the dynamic analysis is to investigate the evolution of strategies *towards* optimality, we cannot assume that, in every period, consumers' decision rules are optimal with respect to the frequency distribution of tariffs in that period. But the range of logically coherent decision rules that are applicable to the consumer's problem is far too wide for all such rules to be included in a tractable dynamic model. We must therefore work with a restricted set of decision rules.

Recall that a consumer h has two types of information at her disposal. For each firm i , she has a noisy signal r_{hi} of the *ex post* utility of choosing that firm. In addition, for each pair of firms i, j with a common standard, she is correctly informed of the ordinal ranking of the corresponding *ex post* utilities. For the purposes of our dynamic analysis, we consider just two decision rules.

The first of these, which will be called *naïve*, implicitly assumes that the fact that two firms use a common standard is not in itself informative (either

positively or negatively) about the prices those firms charge. A consumer who follows the naïve rule begins by comparing the n signals; she provisionally selects the firm, say i , with the highest signal. She then considers the set of firms which use the same standard as i . If this set is a singleton, she chooses firm i ; otherwise, she chooses the firm in this set that has the highest ordinal ranking. Notice that if all consumers follow this rule and if all firms in the market set the same price p^* , each firm's expected sales are N/n , irrespective of the standards they use. This property can be interpreted as neutrality between firms using individuated standards and firms using common standards.⁹ The naïve rule is consistent with the consumer behaviour that characterises IS equilibrium.

The second decision rule, which we call *savvy*, has a positive bias towards common standards: it implicitly assumes that common standards are indicative of low prices, and that this price advantage outweighs any differences in idiosyncratic value. A consumer who follows this rule begins by asking whether any firms have common standards. If there are any such firms, all firms with individuated standards are eliminated; she then applies the naïve decision rule to the remaining firms.¹⁰ Notice that if every firm in the sample has its own standard, she chooses the firm with the highest signal. This rule is consistent both with the consumer behaviour that characterises IS equilibrium and with that characterising CS equilibrium.¹¹

We do not claim that a bias towards shared standards is intrinsically 'rational'; our interest is in whether such a bias might evolve in repeated interactions

⁹ Contrast the rule which first uses the ordinal information to eliminate tariffs that are clearly inferior to others in the sample, and then chooses the tariff which, in the set of non-eliminated tariffs, has the highest signal. If all consumers use this rule and if all firms set the same price p^* , firms with individuated standards have higher expected sales than firms with shared standards.

¹⁰ A variant of the *savvy* rule favours standards that are shared by more firms to those that are shared by fewer. Thus, in a sample of six firms of which three share standard s' and two share s'' , the variant rule chooses the firm which, of those using s' , has the highest ordinal ranking. Had we used this variant in our analysis, our main conclusions would have been unaffected. The most significant effect of using the variant rule is to eliminate equilibria in which there are two or more common standards.

¹¹ It might seem natural to add a third decision rule to the model – the rule that is the mirror-image of 'savvy', favouring individuated standards over shared ones. But having three rules rather than two makes the dynamic analysis much more complicated, while having little effect on the main properties of the model. To see why this is so, notice that the mirror-image rule is the best of the three for consumers, only in a case in which the price set by firms with individuated standards is lower than the price set by firms with common standards. We will show that, in the only conditions under which this can be the case, common-standard firms are more profitable than individuated-standard firms. Thus, the profit-seeking tendencies of firms work to eliminate the conditions under which the mirror-image rule would be beneficial to consumers.

when individuals learn by experience. However, it is interesting to note that a heuristic with some similarities to the savvy rule has been observed by psychologists in situations in which it is apparently irrational. The heuristic of *asymmetric dominance* applies to decision problems in which three options x , y and z are located in two dimensions of value; y dominates z while x neither dominates nor is dominated by either y or z . It seems that individuals treat the dominance relation between y and z as a positive indicator of the value of y relative to x (Shafir and Tversky, 1993). Analogously, if there are three tariffs i , j and k , and if j is known to be better than k , the savvy rule treats this as a positive indicator of the value of j relative to i . This analogy suggests that the savvy rule may have some psychological plausibility, independently of any optimality properties.

The proportion of savvy consumers is denoted by z . This proportion may differ between periods, but in any given period it is treated as parametric. (To avoid clutter, we do not use any notation to identify specific periods.)

Fixed-standard and randomised-standard sectors. We assume that, in any given period, the set of firms can be partitioned into two sectors (either of which may be empty). Firms in the *randomised-standard (RS)* sector set their standards at random (as all firms do in IS equilibrium). All firms in the *fixed-standard (FS)* sector use a fixed common standard s^* (as all firms do in CS equilibrium). The number of FS firms is n^F ; this may differ between periods. For each period, we take the values of n^F and z as given, and define a *temporary equilibrium* as an n -tuple of firms' prices, such that no firm can increase its profit by changing its price, given its own standard, the prices and standards of other firms, and the decision rules used by consumers. Our dynamic analysis focuses on the evolution of n^F and z over time.

How restrictive are these assumptions about standards? First notice that, given the general structure of the model, there is no reason for a firm to use an individuated but non-randomised standard. Given that a firm is using an individuated standard, its demand is unaffected by whether that standard was chosen deterministically or selected at random. However, a firm (say j) which

intends to set an individualized standard puts itself at a strategic disadvantage if it fails to randomise, since if another firm k uses the same standard but sets a marginally lower price, k captures the whole of j 's intended sales (in addition to the sales k would have made by randomisation).

The assumption that there is only one common standard is a convenient simplification, but is not essential for our main results. Since we are analysing the limiting case as $\beta \rightarrow 0$, any two or more firms which use a common standard are effectively in Bertrand competition with one another (whatever the mix of naïve and savvy consumers); in temporary equilibrium, the price set by these firms must be infinitesimally close to marginal cost. If there were two common standards, each used by two or more firms, the prices associated with those standards would be infinitesimally close to one another in temporary equilibrium.¹² Thus, all that really matters for the analysis is the total number of firms which use common standards (that is, standards that are common to two or more firms); whether there is one or more such standard is immaterial.

Our definition of temporary equilibrium implicitly assumes that changes in the distribution of firms between RS and FS sectors take place more slowly than changes in prices. That assumption might be justified on the grounds that it is usually easier for firms to adjust their prices than to adjust their standards. The latter claim may seem more realistic for some real-world markets than for others; and it might be argued that developments in information technology are allowing firms to change their standards much more quickly and at much less cost than in the past. For example, information technology has presumably reduced the administrative cost to a firm of changing the structure of multi-part tariffs; reorganising the layout of a website is surely less costly than reorganising that of a supermarket. However, a transition between the RS and FS sectors involves more than a change in the firm's interface with consumers. It also involves a step change in the level of production and

¹² Suppose that, in temporary equilibrium, firms using some common standard s' set a marginal-cost price p' , while firms using another common standard s'' set a marginal-cost price p'' . The supposition $p' \neq p''$ implies a contradiction, since the firms with the lower price must have both lower sales (because their marginal cost is lower, and marginal cost is non-decreasing) and higher sales (because of the decision rules used by consumers).

sales: typically, FS firms sell more than RS firms. It seems reasonable to assume that changes of this kind take place over a longer time scale than incremental changes in prices within either sector.

Symmetric temporary equilibrium. We assume that, for each (z, n^F) pair, there is a unique temporary equilibrium in which all RS firms set the same price, which we denote p^R . (There may also be non-symmetric temporary equilibria, but our analysis is concerned only with symmetric equilibria.)

Having set out our simplifying assumptions, we proceed to the dynamic analysis itself.

The nature of the dynamics is very different depending on whether n^F is greater than 2. To see why, consider the case $n^F = 2$. In this case, each FS firm has the unilateral power to eliminate the common standard by randomising its own standard. In effect, each FS firm has the power to shift the market to IS equilibrium. Since profits are strictly positive in IS equilibrium, while firms in Bertrand competition make zero profit, we should expect this power to be used. For the reasons we presented in defending the assumption of ‘fixed-standard and randomised-standard sectors’, no firm would want to be the single FS firm in the case $n^F = 1$: if one of two FS firms switches to randomised standards, the other will follow. Thus, if the value of n^F falls below 3, there will be a collapse to IS equilibrium. In contrast, if $n^F \geq 3$, each firm has to take the existence of the common standard as given; it can choose only whether or not to use that standard.

We now consider the dynamics of the model with $n^F \geq 3$. The details of the analysis are explained in the Appendix. Here, we restrict ourselves to the main results and the intuitions behind them. Consider any (z, n^F) such that $n > n^F \geq 3$ and $1 > z > 0$. Taking these values as given, we can define a symmetric temporary equilibrium. Let p^F , q^F and π^F be respectively the price, quantity sold and profit of each FS firm, and let p^R , q^R and π^R be the corresponding values for each RS firm. We assume that if $\pi^F > \pi^R$, there is a

tendency for firms to move from the RS sector to the FS sector and hence for n^F to increase, and conversely if $\pi^F < \pi^R$. If $p^F < p^R$, the expected payoff for consumers is greater if they use the savvy rule rather than the naïve rule; so we assume that in this case there is a tendency for z to increase. Conversely, if $p^F > p^R$, there is a tendency for z to fall.

Figures 1 and 2 (see pages 35 and 36) are phase diagrams showing two possible configurations of the dynamics. As we show in the Appendix, $p^F \geq p^R$ implies $\pi^F > \pi^R$. Thus, the boundary conditions $\pi^F = \pi^R$ and $p^F = p^R$ divide (z, n^F) space into at most three regions: *region A* in which $\pi^F < \pi^R$ and $p^F < p^R$, *region B* in which $\pi^F > \pi^R$ and $p^F < p^R$, and *region C* in which $\pi^F > \pi^R$ and $p^F > p^R$.

In every temporary equilibrium, RS firms set price above marginal cost. Irrespective of the values of p^F and p^R , these firms sell only to naïve consumers. FS firms are in Bertrand competition and price at marginal cost. At low values of z , the high proportion of naïve consumers allows RS firms to charge more than the competitive price c and sell sufficiently large quantities to make positive profit. FS firms produce above MES; pricing at marginal cost, they set $p^F = c$ and make zero profit. This is region A.

At higher values of z , the profit-maximising price for RS firms remains greater than the competitive price but, because there are relatively few naïve consumers, profits for RS firms are negative. Again, FS firms produce above MES, set $p^F = c$ and make zero profit. This is region B.

At the highest values of z , it is possible that the demand faced by RS firms is so low that p^R , despite being above marginal cost, is *less* than the competitive price. (Recall that marginal cost may be increasing below MES.) As in region B, profits for RS firms are negative. But if RS firms are producing below the MES level, FS firms must be producing above it; so yet again, FS firms set $p^F = c$ and make zero profit. This is region C. Whether there is such a region depends on the value of $C'(0)$, i.e. marginal cost at zero output, and on the

price elasticity of conjectural demand for RS firms. Figure 1 shows the case in which there is a region C, while Figure 2 shows the case in which there is not.

The arrows at the top edges of the two diagrams indicate a special property of the dynamics at $n^F = n$. At $n^F = n$, all firms use common standards, and so the distinctive feature of the savvy rule, namely its rejection of individuated standards, has no bite. Thus, the naïve and savvy decision rules give the same expected payoff, and there is no tendency for change in the value of z .

Let E be the set of (z, n^F) points at which $n^F = n$, and with the property that any firm which unilaterally deviated to the RS sector would earn negative profit (while FS firms earn zero profit). Notice that E is non-empty in both diagrams: it is the set of points at which $n^F = n$ and $z > z'$. At each point in E , the following is true: all firms use a common standard and charge the competitive price; all consumers (savvy and naïve) follow decision rules which respect ordinal information; and the proportion of consumers whose (savvy) decision rules reject individuated standards is sufficiently large that no firm can increase profit by deviating from the common standard. Thus, each of these points is a CS equilibrium. Starting from any such point, if there are small perturbations of z or n^F , the dynamics will restore the system to some point in E . In this sense, CS equilibrium is locally stable. Further, the basin of attraction of the set E includes all (z, n^F) such that $n^F \geq 3$ and z is sufficiently close to 1. Of course, we have not yet *proved* that CS equilibrium has these properties; we have merely *illustrated* these properties in two conjectured configurations of the dynamics. The proofs are given in the Appendix. In the Appendix, we also explain the conjectures that lie behind certain features of the diagrams (such as the existence of region A, and the downward-sloping boundary between regions A and B) that are not essential for our main argument.

Finally, we note one further feature of the dynamics. Suppose an additional constraint is imposed on the model, such that at least two specific firms are

constrained to use the fixed standard s^* , irrespective of the profits earned in the two sectors. Thus, the value of n^F cannot fall below 2, and when $n^F = 2$, neither of the FS firms has the power to eliminate the common standard. It is easy to see from the dynamic configurations of Figures 1 and 2 that, if this constraint is imposed, all feasible (z, n^F) points are in the basin of attraction of the set of CS equilibria.

5. Interpretation

In respect of our formal model, the most significant conclusion is that both IS and CS equilibria exist and are locally stable. This raises the question of how a transition from one type of equilibrium to the other might occur. More specifically, are there reasons to expect common standards to emerge spontaneously? Failing that, what kinds of policy interventions might facilitate the evolution of CS equilibrium?

If we remain within the confines of the model, we can answer these questions only in terms of basins of attraction. In order for the dynamics of the model to lead to CS equilibrium, it is first necessary to be at a point at which at least three firms are using a common standard, and at which the proportion of savvy consumers is sufficiently high to ensure that, given the existence of the common-standard firms, individuated-standard firms are unable to avoid making losses. If we want to ask whether such conditions are likely to occur in reality, we must step outside the model. Are there real-world mechanisms which could induce these conditions, so that evolution towards CS equilibrium could then take hold?

On the consumer side, we are looking for mechanisms which tend to sustain high values of z – that is, to sustain the strategy of favouring firms which use common standards – even in markets in which, and at times at which, common standards are rarely observed. The stronger the background propensity of consumers to favour common standards, the more likely it is that a transitory episode in which such standards are used will initiate a self-reinforcing process of movement to CS equilibrium. One mechanism that

might have this effect works through the *generality* of the rule of ‘favouring common standards’. The core idea is to favour tariffs or products which, by virtue of meeting common standards, facilitate value-for-money comparisons with their rivals. This rule is not tied to any specific standard, to any specific firm or firms, or to any specific type of product. Since, in any given market, either IS or CS equilibrium can be sustained, we might expect consumers’ experience *of markets in general* to include instances both of common standards and of individuated ones, and to provide evidence of the association between common standards and low prices. Thus, the rule might be learned in one context and then applied in others.¹³

As far as firms are concerned, we need to ask the following question: starting from a situation in which firms’ standards are individuated, are there dynamic or stochastic processes that might induce episodes in which a small number of firms temporarily use common standards? In thinking about this issue, we should recognise that 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 (see Section 4). 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

¹³ 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).

and non-competitive prices can persist only in combination with episodes of undercutting. In other words: in a realistic form of IS equilibrium, individualized standards will be the norm, but there will be occasional episodes of price competition between firms which are temporarily using common standards. To the extent that price competition, when it does occur, is associated with common standards, consumers are rewarded for using the rule of favouring such standards. This provides an additional reason for expecting that, in a realistic form of IS equilibrium, consumers might learn to favour common standards.

The tendency for episodic price competition will be greater if, contrary to another of the simplifying assumptions of our formal model, different standards are not completely symmetrical with one another. If consumers have preferences between standards, or if some standards are more costly for firms than others, an IS equilibrium will be a state of affairs in which some firms make more profit than others, by virtue of using 'better' (that is, more preferred or less costly) standards. We should expect some form of strategic competition between firms seeking to position themselves at the better standards. This seems likely to produce occasional periods of price competition at those standards.

The implication of all this, we suggest, is that 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. At least some consumers will learn to associate common standards with relatively low prices. The more consumers favour common standards, the greater the incentive for firms to undercut one another. (The firm that undercuts not only takes sales from the firm that is undercut, but also attracts savvy consumers.) A state of turbulence with these general characteristics seems capable of inducing, if only infrequently, the preconditions for evolution to CS equilibrium.

Of course, we must also consider the possibility of transitions from CS to IS equilibrium. As Figures 1 and 2 (pages 35 and 36) show, the main threat to an existing CS equilibrium is a fall in the proportion of savvy consumers. If *all* firms use a common standard, naïve consumers incur no penalty, and so the proportion of savvy consumers can drift downwards. Depending on the precise configuration of the dynamics, such drift might take us into the basin of attraction of IS equilibrium.¹⁴ Thus, the long-run stability of CS equilibrium requires some mechanism that continues to reward the use of the savvy rule. In abstract theoretical terms, random perturbations in the value of n^F would supply such a mechanism, inducing a long-run tendency for movement to (z'', n) in Figure 1 or to $(1, n)$ in Figure 2. In more realistic terms, what is required is some non-zero probability that, at any time, the market contains maverick firms which set non-standard tariffs whose prices are systematically above the competitive level. Again relaxing some of the simplifying assumptions of the formal model, this could be explained in terms of heterogeneity in the population of consumers. Suppose, for example, that a small proportion of consumers have intrinsic preferences for particular standards. Such a standard, if not also the common standard of the CS equilibrium, might provide a niche for a high-price, individuated-standard firm.

Our tentative conclusion is that there may be general market mechanisms which, in the long run, favour the evolution of common standards. 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 – for example, requiring staple food products to be labelled with prices expressed in terms of a stipulated unit of quantity, or requiring all offers of loans to express interest rates in a standard ‘annual percentage rate’ form. 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, the regulation is effectively self-enforcing. Further, it may not be necessary to impose the regulation on

¹⁴ Notice that not all points in region A are necessarily in the basin of attraction of IS equilibrium. Starting from points in A but close to the boundary with B, it is possible that the dynamics lead into region B, and then to CS equilibrium.

all firms; all that is needed is that the number of firms that are required to use a common standard is enough to initiate a self-reinforcing process of transition. Under the assumptions of our model, the regulation of two firms is enough (see the final paragraph of Section 4).

Nevertheless, it is important to understand the self-regulating powers of the market system. We should be cautious about inferring, from the growing evidence of the cognitive limitations of economic agents, that when markets offer 'hundreds of thousands of options', that is too many. We need to take account of how conventions might evolve to help boundedly rational consumers navigate the complexity of the market. Perhaps economics can learn something from Gingrich's Wal-Mart customer.

Appendix: Derivation of the dynamic properties of the model

Proof of claims made in Section 4

Let E be the set of (z, n^F) points such that $n^F = n$ and any firm that unilaterally deviated to the RS sector would earn negative profit. We now show that this set of CS equilibria is non-empty, and that its basin of attraction includes all (z, n^F) such that $n^F \geq 3$ and z is sufficiently close to 1. Our proof considers the temporary equilibrium associated with any given (z, n^F) where $n > n^F \geq 3$ and $1 > z > 0$. We use $p^F, q^F, \pi^F, p^R, q^R$ and π^R to denote the temporary equilibrium values of price, quantity and profit for FS and RS firms.

Step 1. We show that $p^F \geq p^R$ implies $\pi^F > \pi^R$. Since FS firms are in Bertrand competition, the price-cost margin for such firms is zero, i.e. $p^F = C(q^F)$. RS firms sell only to naïve consumers. Because naïve consumers use noisy signals to differentiate between tariffs, each RS firm faces a downward-sloping conjectural demand curve. Thus, price-cost margins for RS firms are positive, i.e. $p^R > C(q^R)$. Now suppose $p^F \geq p^R$. By virtue of the properties of the price-cost margins, this implies $C(q^F) > C(q^R)$. Since marginal cost is non-decreasing, and since $N/n \geq q^*$, this is possible only if $q^F > q^* > q^R$. The inequality $q^F > q^*$ implies $p^F = C(q^F) = c$ and $\pi^F = 0$. Since average cost is decreasing at quantities less than q^* , $q^R < q^*$ implies $C(q^R)/q^R > c$. But $p^R \leq p^F = c$. Thus, RS firms are selling at a price below average cost, and so $\pi^R < \pi^F$.

Step 2. The result from Step 1 establishes that, as claimed in the main text, the boundary conditions $\pi^F = \pi^R$ and $p^F = p^R$ divide (z, n^F) space into at most three regions: region A in which $\pi^F < \pi^R$ and $p^F < p^R$, region B in which $\pi^F > \pi^R$ and $p^F < p^R$, and region C in which $\pi^F > \pi^R$ and $p^F > p^R$.

Step 3. Given the continuity properties of the model, the result from Step 1 also establishes that regions A and C, if they exist, must be separated by region B. (If A and C shared a common boundary, points on that boundary would have $\pi^F = \pi^R$ and $p^F = p^R$.)

Step 4. If $p^F < p^R$, it follows from the specification of the consumers' decision rules that $q^F > q^R$, irrespective of the value of z . Thus, since $N/n \geq q^*$, we have $q^F > q^*$. Since FS firms price at marginal cost, this implies $p^F = c$ and $\pi^F = 0$. It has been shown as part of Step 1 that, if $p^F \geq p^R$, then $p^F = c$ and $\pi^F = 0$. Thus $p^F = c$ and $\pi^F = 0$ everywhere.

Step 5. We now show that region C cannot include values of z arbitrarily close to 0. Suppose, to the contrary, that there is a temporary equilibrium with $p^F > p^R$ and $z \approx 0$. Since $p^F > p^R$, each RS firm must sell to more naïve consumers than each FS firm does; since almost all consumers are naïve, we have $q^F < q^R$. But from Step 1, $p^F > p^R$ implies $q^F > q^R$, a contradiction.

Step 6. Notice that, for all n^F such that $3 \leq n^F < n$, $q^R \rightarrow 0$ as $z \rightarrow 1$. (This follows immediately from the fact that RS firms sell only to naïve consumers.) Since fixed costs are strictly positive, $\pi^R < 0$ as $z \rightarrow 1$, while $\pi^F = 0$. Thus, points at which $z \approx 1$ are either in region B or in region C (or on the boundary between these regions).

Step 7. We now show that, starting from any (z, n^F) point at which $p^F = p^R$, increases or decreases in n^F (with z constant) do not affect the temporary equilibrium values of $p^R - p^F$ or $\pi^R - \pi^F$. Consider any RS firm j . It sells to, and only to, those naïve consumers for whom its signal is the highest. The expected number of such consumers depends only on the total number of naïve consumers and the price charged by each firm. Thus, if $p^F = p^R$, the conjectural demand function faced by any RS firm j is independent of the distribution of other firms between the RS and FS sectors. So changes in n^F do not affect the temporary equilibrium values of p^R and π^R . But, from Step 4, $p^F = c$ and $\pi^F = 0$ everywhere. So changes in n^F do not affect $p^R - p^F$ or $\pi^R - \pi^F$. This implies that if region C exists, the boundary between it and region B (i.e. the locus of points at which $p^R = p^F$) is a vertical line in (z, n^F) space. (Conceivably, there could be a 'thick' boundary zone with vertical edges. For simplicity, we do not consider this case, but our proof can be extended to take account of it.)

Step 8. Suppose region C exists. Then, from Steps 3, 5, 6 and 7, there is some z' such that $0 < z' < 1$, all points to the right of the line $z = z'$ are in Region C, and points immediately to the left of the line are in Region B. Then there exists a set E of CS equilibria which contains (at least) all points (z, n^F) such that $n^F = n$ and $z \geq z'$; and the basin of attraction for E includes the whole of region C, and *a fortiori* all (z, n^F) such that $n^F \geq 3$ and z is sufficiently close to 1 (see Figure 1).

Step 9. Suppose region C does not exist. Then, from Step 6, there is a region B which includes all (z, n^F) such that $n^F \geq 3$ and z is sufficiently close to 1. For some $z' < 1$, the set E of points at which $n^F = n$ and $z > z'$ is a set of CS equilibria, and the basin of attraction for E contains the whole of region B (see Figure 2).

Steps 8 and 9 together establish that, for some $z' < 1$, the set E of points at which $n^F = n$ and $z > z'$ is a set of CS equilibria, and that basin of attraction for E includes all points such that $n^F \geq 3$ and z is sufficiently close to 1.

Further properties of the dynamics

As we have drawn them, Figures 1 and 2 (pages 35 and 36) have the following three common properties in addition to those derived above: region A exists; it includes all (z, n^F) such that $n > n^F \geq 3$ and z is sufficiently close to 0; and the boundary between regions A and B is downward-sloping. These properties are not essential to our analysis. We conjecture that they hold for most plausible specifications of the functions in the model, but we cannot provide complete proofs that they do. We now explain these conjectures.

First, consider the effects on the relative profitability of FS and RS firms of an increase in z with n^F constant. We know from Step 4 of the proof that this increase in z will have no effect on p^F (which will remain equal to c) or on π^F (which will remain equal to zero). Since RS firms sell only to naïve consumers, it seems reasonable to expect that a reduction in the proportion of

such consumers, associated with no change in the number of RS firms or the price charged by FS firms, will reduce the value of π^R , and hence that of $\pi^R - \pi^F$.

Next, consider the effects on the relative profitability of FS and RS firms of an increase in n^F with z constant, starting from a point at which $p^F < p^R$ (i.e. from a point in region A or region B, or on the boundary between these regions). We know from Step 4 that this increase in n^F will have no effect on p^F (which will remain equal to c) or on π^F (which will remain equal to zero). For any firm j which remain RS, q_j depends only on the total number of naïve consumers and the price charged by each firm; the distribution of other firms between the RS and FS sectors influences q_j only through its effect on the distribution of prices (compare Step 7). The effect on j of an increase in n^F is to increase the number of its competitors charging the relatively low price p^F . Intuitively, it seems reasonable to expect that this effect, associated with no change in the number of naïve consumers, will reduce the value of π^R , and hence that of $\pi^R - \pi^F$.

Suppose the conjectures of the two previous paragraphs are true. Then, if region A exists, the boundary between it and region B is a downward-sloping curve, with A to the left. Similarly, if region C exists, the boundary between it and region B is a vertical line, with C to the right. We know from Step 5 that C does not include values of z that are arbitrarily close to zero. So at values of z sufficiently close to zero, we cannot be either in or on the boundary of region C, i.e. we must have $p^F < p^R$. By Step 4, $p^F = c$ and $\pi^F = 0$ everywhere. Let π'_j be the profit that would be earned by any RS firm j if it unilaterally deviated from temporary equilibrium by setting $p_j = c$ while still randomising its standard. Since p^R is the price which maximises j 's profit, given the tariffs of other firms, we know that $\pi'_j \leq \pi^R$, and it seems reasonable to expect this inequality to be strict. But j 's sales to naïve consumers depend only on its own price and on the prices charged by other firms; they are independent of whether its standard is randomised. Thus as $z \rightarrow 0$, $\pi'_j \rightarrow \pi^F$. This motivates

the conjecture that $\pi^F < \pi^R$ at sufficiently low values of z . In other words, points at which $z \approx 0$ are in region A.

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.
- Ellison, Glenn (2005). A model of add-on pricing. *Quarterly Journal of Economics* 120: 585-637.
- Ellison, Glenn and Sara Fisher Ellison (2004). Search, obfuscation, and price elasticities on the internet. Discussion paper, MIT and NBER.
- Gabaix, Xavier and David Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121: 505-540.
- 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.
- 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.
- Sunstein, Cass and Richard Thaler (2003a). Libertarian paternalism is not an oxymoron. *University of Chicago Law Review*, 70: 1159-1202.

Sunstein, Cass and Richard Thaler (2003b). Libertarian paternalism.
American Economic Review, Papers and Proceedings, 175-179.

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: One configuration of the dynamics




Figure 2: An alternative configuration

