University of East Anglia

CENTRE FOR COMPETITION POLICY

Transactional fairness and unfair price discrimination in consumer markets

Bruce Lyons
Centre for Competition Policy & School of Economics
University of East Anglia

Robert Sugden
Centre for Competition Policy & School of Economics
University of East Anglia

CCP Working Paper 20-07
This version: October 2020

There is growing public concern about the 'unfairness' of many pricing practices that have become common in consumer, particularly digital, markets (e.g. auto-renewal at a high price, expensive default add-ons). Industrial and behavioural economists have developed theories that explain the conditions under which these practices are profitable for firms, and their implications for consumer welfare. We argue that there is a mismatch between the welfare economic principles on which this theoretical work is grounded and the normative perspective in which the pricing strategies in question are viewed as unfair. As a result, when regulators look to economics for guidance about fair pricing, they struggle to reconcile two fundamentally different normative approaches. We develop a concept of 'transactional fairness', grounded in the normative approach of Sugden's 'Community of Advantage', that is reflective of public concerns. Transactional fairness requires satisfaction of 'no deception', 'no hindrance' and 'public explanation' criteria. It is complementary to established welfare criteria of economic efficiency and distributional equity, but is based entirely on the relationship between individual buyer and seller. Transactional fairness establishes clear principles with realistic information requirements that are appropriate for compliance by firms. The approach potentially helps restore public faith in markets without either deterring the emergence of (non-deceptive and non-hindering) business models, or requiring frequent ad hoc fire-fighting interventions by regulators.

Contact Details:
Bruce Lyons: b.lyons@uea.ac.uk
Robert Sugden: r.sugden@uea.ac.uk

ISSN 1745-9648

Transactional fairness and unfair price discrimination in consumer markets

Bruce Lyons* and Robert Sugden**

31 August 2020

* Centre for Competition Policy and School of Economics, University of East Anglia, Norwich, NR4 7TJ, UK

Email: b.lyons@uea.ac.uk

** Centre for Behavioural and Experimental Social Science, Centre for Competition Policy and School of Economics, University of East Anglia, Norwich, NR4 7TJ, UK

Email: r.sugden@uea.ac.uk

Abstract

There is growing public concern about the 'unfairness' of many pricing practices that have become common in consumer, particularly digital, markets (e.g. auto-renewal at a high price, expensive default add-ons). Industrial and behavioural economists have developed theories that explain the conditions under which these practices are profitable for firms, and their implications for consumer welfare. We argue that there is a mismatch between the welfare economic principles on which this theoretical work is grounded and the normative perspective in which the pricing strategies in question are viewed as unfair. As a result, when regulators look to economics for guidance about fair pricing, they struggle to reconcile two fundamentally different normative approaches. We develop a concept of 'transactional fairness', grounded in the normative approach of Sugden's 'Community of Advantage', that is reflective of public concerns. Transactional fairness requires satisfaction of 'no deception', 'no hindrance' and 'public explanation' criteria. It is complementary to established welfare criteria of economic efficiency and distributional equity, but is based entirely on the relationship between individual buyer and seller. Transactional fairness establishes clear principles with realistic information requirements that are appropriate for compliance by firms. The approach potentially helps restore public faith in markets without either deterring the emergence of (non-deceptive and non-hindering) business models, or requiring frequent ad hoc fire-fighting interventions by regulators.

Keywords: Price discrimination, unfair pricing, consumer law, competition policy

JEL Codes: D61, D63, K21, K23, L40, L51

Funding: This research was supported by the Economic and Social Research Council, grant number ES/P008976/1. Sugden's work was also partly funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement No. 670103.

There is growing public concern, expressed in the media, in political debate and by consumer advocacy groups, about what is thought to be the 'unfairness' of many forms of price discrimination that have become common in consumer markets. Undoubtedly, current economic and technological trends are facilitating new forms of price discrimination by firms. Among these trends are the increasing importance of personalised and digital services in the economy; the shift in selling strategies from sale to rental and from payment-per-item to subscription payment; the shift to direct debit payment and auto-renewal for ongoing service contracts; and developments in information technology that give firms access to data about the individual characteristics of their customers and allow firms to use this information to set individuated prices. Public concern about price discrimination is understandable.

Regulatory agencies are very conscious of this concern, and have responded with inquiries into unfair pricing strategies and by introducing regulations designed to limit their use.[1] Industrial and behavioural economists have developed theories that explain the conditions under which various kinds of price discrimination are profitable for firms, their implications for consumer welfare, and the likely effects of different regulatory responses. However, the premise of our paper is that there is a mismatch between the welfare economic principles on which this theoretical work is grounded and the normative perspective in which the pricing strategies in question are viewed as unfair. As a result, when regulators look to economics for guidance about fair pricing, they struggle to reconcile two fundamentally different normative approaches. Our paper is an attempt to resolve this problem by formulating a concept of 'transactional fairness' that can represent the intuitive sense of fairness that is invoked in public debate, and that can be used by regulators and in economic analysis.

We begin with two familiar examples of price discrimination that illustrate the mismatch we have described. Here, and elsewhere in the paper, we follow a conventional

[1] For example, the major UK regulators have recently become very active in trying to understand 'unfair pricing' both in relation to 'vulnerable consumers' and more widely. This includes market investigations that identified apparently 'unfair' pricing practices such as local energy customers being offered higher tariffs than 'out of area' customers (Ofgem, 2009), automatic reversion to an unfavourable 'single variable tariff' at the end of fixed term energy contract (CMA 2016a), 'back-book' bank customers being left on unfavourable accounts that are no longer available to new customers (FCA, 2015; CMA 2016b), bundled contracts for 'mobile phone plus usage' that continue by default after the phone has been paid for so a SIM-only contract would be much cheaper (Ofcom 2019). More general investigations have highlighted how widespread these practices are becoming (CMA, 2018; FCA, 2018b), and have resulted in follow-on investigations and actions (FCA, 2019; Ofcom, 2020).

practice in the economic literature by distinguishing between *savvy* and *naïve* consumers ('savvies' and 'naïves' for short). Savvies are rational in the sense of neoclassical theory: they act on stable, internally consistent and context-independent preferences – for short, preferences that are *integrated* – and on correct beliefs about relevant economic variables, including search opportunities and prevailing business practices. Naïves fall short of these standards in some way, for example by lacking essential information, holding false beliefs, or being susceptible to psychological influences that neoclassical theory assumes away. Many allegedly unfair pricing practices discriminate between savvies and naïves.

Our first example is the practice that leads to *bill shock*. A firm offers a service contract with a below-cost or even zero price for its core component, along with a schedule of charges, well in excess of cost, for add-ons that consumers can incur while using the service and will be billed for afterwards. Familiar examples include late payment fees, bank charges for unarranged overdrafts, and high unit prices for mobile phone usage above some threshold. Savvies avoid the add-ons and benefit from the low headline prices. Naïves incur the add-ons, either by signing up to contracts without reading (or recognising the significance of) the small print or by inattention afterwards.

Our second example is the *loyalty penalty*. Firms offer service contracts that are subject to periodic renewal; at each renewal, the firm quotes a new price. Firms attract new customers by offering low initial prices. The longer a customer has stayed with a firm, the higher the price she pays. A common strategy is to increase a customer's price at each renewal by proportionately more than the increase in the firm's costs, tailoring each increase so that it is not obviously excessive ('price walking'). Savvies notice the price increases and switch frequently between suppliers. Naïves do not notice and are penalised for their 'loyalty' to their original supplier.

In these examples, price discrimination is based on differences in consumers' information about, understanding of, or attention to firms' pricing strategies. Because of this, the textbook result that price discrimination cannot occur in a competitive market does not apply. In a market with no fixed costs or barriers to entry, competition between firms can induce an equilibrium in which firms operate at minimum cost and earn only normal profit, but there is still price discrimination between naïves and savvies. In the bill shock case, firms compete to offer low headline prices for the core service, seeking to attract naïve customers who will incur add-on charges. In the loyalty penalty case, firms compete to offer low-price introductory offers, seeking to attract naïve customers who will continue to renew their

contracts as the price increases. In both cases, savvies are cross-subsidised by naïves (Gabaix and Laibson (2006); Armstrong and Vickers, 2012; Armstrong, 2015; Grubb, 2015).

Such pricing practices are seen by many consumers as unfair. For example, a large-scale telephone survey of UK motor and home insurance customers, commissioned by the Financial Conduct Authority, asked respondents to say whether the following was fair or unfair: 'Mr Smith has been with the same insurance firm for 5 years and pays £500 for his buildings insurance. Mr Jones, whose house is identical, asks Mr Smith's insurance provider for a quotation, and is quoted £300 for the same policy'. Only 9 per cent of respondents thought this 'fair' (FCA, 2019).[2] Intuitively, it is easy to understand how this kind of discrimination can be seen as unfair, but (as we will explain in more detail later) it is more difficult to represent its unfairness within the theoretical framework of normative economics.

One might ask why this is a problem for economics. If economics has its own coherent methods of normative analysis, and if, according to an analysis conducted using those methods, there is nothing objectionable about some pricing strategy that the general public regards as unfair, isn't the problem simply that the general public has insufficient understanding of economics? We believe that such a response would be inadequate, for (at least) the following reasons.

First, we conjecture that many professional economists – including specialists in welfare economics and industrial organisation – will think that the practices we have described *are* unfair. If we are right about this, the problem must surely be more than a misunderstanding of economics.

Second, much work in economics is framed as leading to 'policy implications' – as giving advice to market regulators and economic policy-makers. In a democratic society, regulators and policy-makers are not benevolent autocrats, free to act on advice from whoever they choose to consult; ultimately they are constrained by, and so need to take account of, citizens' judgements about economic matters. An essential first step in doing this is to understand and characterise those judgements. We believe there should be a default presumption that widely held judgements about economic fairness and unfairness are

[2] This was not because respondents were unaware of pricing strategies in the insurance market: 89 per cent of respondents agreed with the statement 'Typically, first time customers receive a lower price'. Nor did respondents think it unfair that search was rewarded: 80 per cent thought it fair that 'Alex gets her insurance renewal letter. She shops around using a price comparison website and gets an offer from a different insurance provider and saves £75.'

intelligible within some reasonably coherent normative framework that many citizens implicitly endorse. Identifying such a framework, if it exists, should be seen as an issue on the agenda of economics.

Third, if market regulation is to be effective, it has to be based on general principles that are stated publicly and are commonly understood by actors in the economy: it should not be an unsystematic collection of case-by-case ex post decisions. In the context of price discrimination, firms need to know in advance the principles by which regulators determine which practices are permissible, and consumers need to know in advance what they can and cannot expect of firms. If regulation will in fact be influenced by citizens' conceptions of fairness, those conceptions need to be codified.

Fourth, if a market economy is to be politically sustainable, it is important that consumers have a general sense that the market treats them fairly. Since most consumers' main contact with the market is through firms, it is important that there is a publicly accepted standard of fair business behaviour to which firms can be seen to comply, or be held to account for violating. If a firm is able to continue to use pricing practices that are generally perceived as unfair, this is liable to undermine trust, not only in that particular firm, but also more generally in markets, regulators and the market system. Effects of this kind are particularly dangerous in markets such as financial services, where trust is fundamental to the nature of the product.

Our final reason goes deeper. The normative assumptions used in current economic analysis are not immutable. Philosophically and methodologically, they reflect the influences of utilitarianism and rational choice theory on modern welfare economics. They support a view of economic institutions as mechanisms through which unintended social outcomes are generated through the interactions of rationally self-interested individuals, and a view of public policy-making as a problem of mechanism design whose objective is the maximisation of social welfare. But other ways of understanding economic life have equally deep intellectual roots. Adam Smith's (1776/ 1976: 456) metaphor of the 'invisible hand' that leads self-interested merchants to promote the interests of society is often (and not unreasonably) cited as a precursor of the First Fundamental Theorem of welfare economics. But in another famous passage, Smith characterises the market in terms of individual freedom – as 'the obvious and simple system of natural liberty' in which every man 'is left perfectly free to pursue his own interest in his own way' (p. 687). Ideas of fairness that are currently

seen as non-economic might be more compatible with an approach to normative economics that emphasises freedom and opportunity.

Our objective is to characterise a concept of fairness, grounded in normative principles of consumer freedom and opportunity, that can represent intuitions that underlie public concerns about many forms of 'unfair' price discrimination. We must emphasise that we are not proposing this concept as a standard that can replace the conventional economic standards of efficiency and distributional equality: we are presenting it as complementary with those standards – as a significant and distinct normative dimension.

1. Standard approaches in neoclassical and behavioural economics

In both neoclassical and behavioural economics, the standard normative approach to economic appraisal uses the yardstick of *economic efficiency* or *social welfare*. In principle, this involves identifying all those affected by the outcome of a policy or business practice, and assessing the impact on each individual. Given the assumption that each individual has integrated preferences over relevant outcomes, these impacts can be measured in monetary units as compensating or equivalent variations. The net sum of these impacts is the effect of the policy or practice on economic efficiency. The same sum is often also treated as an approximate measure of the effect on social welfare. Alternatively, the effect on social welfare is measured as a weighted sum of the effects on individuals, with weights that may be different for different groups of people, classified (for example) by income or need. Such weights can be interpreted as publicly approved normative judgements, or as proxies for the marginal utility of income, interpreted in classical utilitarian terms. Market regulation is often based on a simpler *consumer welfare* standard, which sums unweighted effects on consumers but gives zero weight to effects on corporate profits. This approach may be justified on the grounds that the primary role of regulatory agencies is to ensure that markets work well for consumers; it is the responsibility of firms to work out how to make profit within the rules set by regulators and legislators.[3]

The assumption that consumers have integrated preferences can lead to difficulties in assessing the welfare of naïve consumers. In mainstream economic theory, an individual's preferences are assumed (or, in some versions of the theory, defined) to be revealed in her

[3] This is reflected in the UK Competition and Markets Authority's primary duty 'to promote competition… for the benefit of consumers' (Enterprise and Regulatory Reform Act 2013). The UK sector regulators have similarly consumer-oriented duties.

decision-making behaviour, and the consistency properties that are attributed to those preferences are interpreted as principles of rationality. By definition, however, a naïve consumer is one whose decisions are *not* based on complete information, or are *not* fully rational. In the case of incomplete information, there is a long-standing practice in the industrial organisation literature of using models in which naïves have integrated preferences that would be revealed in decisions made with complete information; what counts as 'complete' information is specified by the nature of the model (e.g. Diamond, 1971; Salop and Stiglitz, 1977; Varian, 1980). In normative versions of behavioural economics, it is common to extend this practice by assuming that a naïve individual has integrated *latent* (or 'true') preferences that would be revealed in her decisions if she had complete information and was not subject to 'biases' or 'errors' attributable to deficiencies of cognitive capacity, attention or self-control (e.g. Sunstein and Thaler, 2003; Camerer et al., 2003; Kőszegi and Rabin, 2007; Bershears et al., 2008).[4] A related approach infers preferences solely from individual choice behaviour, but ignores data from choices made in situations in which psychological effects induce 'incorrect perceptions' of the choice set (Bernheim, 2016).

It would be in accord with these analyses of naïve consumer behaviour to define a firm's behaviour as unfair if it activates psychological biases in a way that induces its customers to act contrary to their true preferences. This idea is central to the EU's 2005 Unfair Commercial Practices Directive.[5] According to this Directive, one condition that can make a commercial practice unfair is that the practice 'materially distorts' the behaviour of an 'average consumer' in relation to the product of the relevant firm; 'material distortion' is defined as 'using a commercial practice to appreciably impair the consumer's ability to make an informed decision, thereby causing the consumer to take a transactional decision that he would not have taken otherwise' (European Union, 2005: Articles 2, 5.2).[6] This condition has some similarities with a condition that will feature in our proposed definition of transactional fairness. But notice the implicit assumption in the 'thereby …' clause that it is possible to predict the decisions that consumers would have taken in the absence of some

[4] Camerer et al. (2003: 1218) draw an explicit analogy between behavioural economics and the economics of information. They argue that from the 1930s, economic theory was successfully extended by relaxing first the assumption of perfect competition and then the assumption of perfect information. Relaxing the assumption of perfect rationality is 'a logical next step'.

[5] This Directive, including the relevant wording, was translated into UK law by the Consumer Protection from Unfair Trading Regulations 2008.

[6] To be unfair, a practice must also be 'contrary to the requirements of professional diligence' (Article 5.2).

allegedly unfair commercial practice, presumably by identifying settings in which consumers' decision-making abilities are *not* impaired. In other words, the Directive is invoking a model of latent preference and bias, and assuming that latent preferences can be identified by screening out the effects of bias.

However, the claim that naïve choices can be explained by the interaction of integrated latent preferences and psychological biases is questionable. Since this issue will turn out to be significant for our analysis of transactional fairness, we consider it briefly here. Infante et al. (2016) have argued that the concept of latent preference lacks psychological foundations and is explanatorily redundant.[7] We explain this critique in relation to a well known contribution to behavioural consumer theory – Bordalo et al.'s (2013) model of the effects of 'salience' on consumer choice.[8]

This model builds on the behavioural finding that, when valuing any good, individuals tend to give most attention to those attributes on which it stands out relative to the goods with which it is being compared. In Bordalo et al.'s leading example, a consumer chooses between a high-price, high-quality French wine and a low-price, low-quality Australian wine. At 'supermarket' prices of $20 and $10 respectively, the price difference is highly salient; at 'restaurant' prices of $50 and $40, price is less salient and quality is more salient. Bordalo et al. model the latent utility that a consumer derives from a good as a linear function of the quantities of the good's positive and negative attributes (quality and price in the example). A 'rational' consumer maximises latent utility. A 'salient thinker' (i.e. naïve consumer) maximises a function in which the weight attached to each attribute in the latent function is 'distorted' in a way that reflects its relative salience. This model can explain the pattern in the choices of the wine buyer in the example, and a range of related anomalies in consumer choice. But, unless special assumptions are used to calibrate the model, it cannot tell us which (if either) of the supermarket or restaurant preferences is rational and which is distorted.[9] The underlying psychological theory proposes that individuals give more

---

[7] Related critiques of this concept are developed by Berg and Gigerenzer (2010) and Rizzo and Whitman (2020: 46–52).

[8] The following discussion of Bordalo et al.'s model is based on Sugden (2015).

[9] Bordalo et al. specify their model so that a consumer's 'rational' preferences are revealed in tasks which elicit her willingness to pay for a single good, considered in isolation. They do not explain why willingness-to-pay tasks have this special status. Perhaps their thought is that, if such a task is presented in the simplest possible form, there can be no external sources of bias. But the accumulated evidence suggests that when valuation tasks are presented in apparently context-free form, responses are often stochastically unstable and influenced by irrelevant cues (e.g. Ariely et al., 2003).

attention to attributes that are more salient; it does not say – and has no need to say – what is the *correct* distribution of attention between them. The salience-dependence of choice is an empirical concept, but correctness is not.

The approaches to normative economics described above – the approaches that are most commonly used in neoclassical *and* behavioural economics – share three significant features. First, they are *synoptic*: they make normative judgements from the viewpoint of a single social planner or ethical observer, looking at society as if from outside. Second, they are *consequentialist*: normative judgements are about the outcomes that individuals experience, interpreted without reference to the procedures by which those outcomes came about. Third, they (even the approaches used in behavioural economics) are ultimately *based on rational choice theory*: effects on the welfare of individuals are defined in relation to preferences that are (or would be) revealed in rational choice, and the social planner or ethical observer is understood as a rational maximiser of social or consumer welfare. These features impose constraints on the kinds of fairness and unfairness that can be taken into account.

Take the case of bill shock in a competitive market with no fixed costs. In equilibrium, firms produce at minimum average cost and earn only normal profit. The profits that firms earn from naïves are competed away in the loss-leading headline prices from which savvies benefit. Suppose for simplicity that the incremental cost to a firm of supplying the relevant add-on is zero, that naïves do not anticipate the possibility of incurring this add-on, and that attentive customers can evade the add-on at no cost to themselves.[10] Under these conditions, both savvies and naïves have the ex ante perception that they are buying only the core service and paying only the headline price; ex post, it is as if the firm's loss from selling below cost is made up by an arbitrary levy on naïves. Apart from the possible inefficiency that results from overconsumption of the underpriced core service by both types of consumer – an effect that seems orthogonal to concerns about fairness – the effect of the price discrimination is a pure transfer from naïves to savvies. In a standard economic analysis, an evaluation of this effect would be a distributional judgement based on the relative income of savvies and naïves. If naïves are predominantly cash-rich, time-poor consumers who can afford to be inattentive to add-on prices, that transfer might be judged an *increase* in welfare.

Even if naïves are distinguished by their cognitive 'vulnerability', for example, due to age, infirmity or lack of financial understanding, it is not obvious that effects on them should

[10] In fact, many bank customers incur overdraft fees through inattention when they have adequate balances in other accounts: see, e.g., Stango and Zinmann (2014).

be given greater distributional weight than effects on savvies. Suppose that Arthur is an 85-year-old widower with a good index-linked pension but declining mental powers, and that Bella is a 30-year old low-income single mother who takes care when using her bank account. As a result of his poor memory, Arthur incurs high add-on bank charges from which Bella ultimately benefits. Viewed in a consequentialist perspective, the effect is a small transfer of income from someone for whom the marginal utility of income is relatively low to someone for whom it is probably higher. But one might still think that, by taking advantage of his vulnerability, Arthur's bank has treated him unfairly.

To provide context for the loyalty penalty, consider some relevant results from one-period models of price search and price dispersion. Price competition can be effective only if consumers seek out low prices, and consumers would have no reason to compare firms' offers if they already knew that all offers were exactly the same. In realistic models of price competition, some price dispersion persists in equilibrium. The extent of this dispersion depends on consumers' tradeoffs between the costs of searching for low prices and the benefits of finding them. On average, consumers who are more willing to search pay lower prices, but those who search more confer a positive externality of lower prices on those who search less (e.g. Diamond, 1971; Salop and Stiglitz, 1977; Stiglitz, 1979; Varian, 1980). Consider two extensions. First, a regulation that imposes a price cap as a means of reducing price dispersion may, by reducing the incentive to search, increase prices at the lower end of the distribution and so, under some conditions, reduce expected welfare for *all* consumers (Armstrong, Vickers and Zhou, 2009). Second, if firms can identify and discriminate between savvy consumers who search and naïve consumers who do not, then each firm has the incentive to set a low price for savvies and a high price for naïves, which removes the positive search externality but may still reduce average price paid and raise average welfare (Armstrong, 2015).[11]

Next, consider the implications of this result for an analysis of the loyalty penalty. Consumers who are already in the market have a default provider. Its client base may include both savvy and naïve consumers. A 'loyalty penalty' pricing strategy for the firm would be to offer existing clients a high price (but not so high as to induce search by naïves) and offer a lower price to savvy clients (including its own and those of rival firms) who discover it. As

[11] Such price discrimination may implemented, at least approximately, by setting a higher price in shops than on the internet (if that is how the savvy search), or by offering a range of complex tariffs out of which only the savvy can identify the single low price, or by using individualised data on search history to set personalised prices.

compared with firms offering a single price, this may be a pro-competitive strategy and reduce average or even all prices (Corts, 1998; Stole, 2007, sec.3.4). Conceivably, a cap on renewal prices might raise the overall price level sufficiently to increase the welfare of many naïves. However, thinking about the way the firm appears to be 'hiding' offers for which its naïve clients are eligible, one might still agree with the survey respondents who thought that Mr Smith's insurance provider had treated him unfairly.

2. Mutual benefit as an ethical standard

Our analysis of transactional fairness builds on an approach to normative economics that is neither synoptic nor conequentialist, and is not based on rational choice theory – an approach developed by Sugden and co-authors (Sugden, 2004, 2018; McQuillin and Sugden, 2012; Bruni and Sugden, 2013; Infante et al, 2016; Isoni et al., 2020).

As historical background: in the early 2000s, behavioural economists began to recognise the need to develop a form of normative analysis that was consistent with behavioural findings. The now-standard method of defining welfare in terms of latent preferences, described in Section 1, was proposed in two influential manifestos. Sunstein and Thaler (2003) linked this method with the proposal that public policy should be based on 'libertarian paternalism': naïve individuals should be 'nudged' towards behaviour that would satisfy their latent preferences, but without being subjected to paternalistic restrictions on choice. Camerer et al. (2003) proposed 'asymmetric paternalism', based on a (not fully defined) behavioural form of cost-benefit analysis that would take account both of the benefits that paternalistic regulations confer on naïves and the costs they impose on savvies. Contemporaneously, Sugden (2004) proposed a radically different approach.

The key idea in this approach is to focus on individuals' *opportunities* – represented by their choice sets – rather than on how far their preferences are satisfied by the choices they make from those sets. Because each individual's choice set can be defined objectively, without reference to her preferences or psychology, a normative criterion that is defined in terms of properties of individuals' choice sets is not disabled by the findings of behavioural economics. Sugden (2004) formulates an opportunity-based analogue of the Pareto-optimality condition of neoclassical welfare economics. Roughly speaking, this Opportunity Criterion requires that individuals' choice sets are such that every group of individuals has the collective opportunity to make any feasible transaction amongst themselves that they might conceivably find mutually acceptable. The only 'acceptability' assumption is that

individuals attach positive value to the medium of exchange (i.e. prefer to buy at low prices and to sell at high prices).

It can be shown that every competitive equilibrium of an exchange economy – including 'storage' economies in which exchange takes place over a sequence of time periods – satisfies the Opportunity Criterion (Sugden, 2004; McQuillin and Sugden, 2012). McQuillin and Sugden (2012: 630–631) say that a person is 'willing to pay for' a good if (at the moment of decision) she is willing to give up what would induce others to supply it. In this sense, their result shows that a competitive market 'gives each person, rational or irrational, what she wants and is willing to pay for, when she wants it and is willing to pay for it'.

Why might this be viewed as a *desirable* property of the market, even if what an individual wants at the moment of decision is not an implication of integrated preferences? Sugden does not claim that this property can be seen as desirable from the synoptic viewpoint of a social planner, seeking to maximise some measure of social welfare that integrates the welfares of different individuals. The claim is that it can be seen as desirable *by each individual separately*, thinking reflectively about what *she* wants from an economic system while recognising that the properties of that system must be justified to everyone. This normative viewpoint is the 'contractarian perspective' (Sugden, 2018: 29–52). And from the individual's viewpoint, the opportunities provided by the market need not be desired as the means of maximising some measure of her own welfare that integrates the consequences of different transactions: the desirability of each transaction lies in the fact that she desires it at the time and in the context that it takes place.

Viewed in the contractarian perspective, market transactions can be seen as voluntary interactions between individuals who are cooperating for mutual benefit. In this context, behaving ethically is playing one's fair part in a cooperative scheme. Sugden (2018: 256–281) expresses this ethic of cooperation in the following principle, further analysed by Isoni et al. (2020):

> *Principle of Mutual Benefit*. When participating with others in a voluntary interaction, and for as long as others' behaviour in that interaction is consistent with this very principle, behave in such a way that the other participants are able to satisfy normal expectations about the consequences of the interaction for them.

The Principle of Mutual Benefit applies to voluntary interactions in general, but market transactions are a paradigm case.

The concept of *normal expectations* applies to a given class of similar interactions that take place recurrently within some population. Normal expectations are beliefs that: (i) are at least approximately correct as a description of actual behaviour in the population, (ii) most members of the population can reasonably infer from their own experience and information, and (iii) are in fact held by most members of the population. In terms of a definition that is widely used in the theory of social norms, patterns of behaviour that are the objects of normal expectations are *descriptive norms* (Cialdini, Reno and Kallgreeen, 1990; Bicchieri, 2006).

The idea behind the Principle of Mutual Benefit is that if some class of interactions is voluntary (i.e. each episode of interaction takes place only with the consent of all participants) and if there are normal expectations about how people behave within such interactions (the *practice* for that class of interactions), then your choosing to participate in an episode of interaction is a signal that, conditional on the others also choosing to participate: (i) you intend to conform to the practice; (ii) you expect the other participants to conform; (iii) you expect to benefit from the interaction; and (iv) you expect the others to benefit. By virtue of these properties of participation decisions, the existence of a practice is an opportunity for mutual benefit. By conforming to a practice in a specific episode, you play your part in a cooperative activity that involves you and the other participants. At the same time, you are also playing your part in a wider scheme of cooperation: by conforming to a practice, you reinforce the common expectation that it will be followed, and so help to sustain opportunities for others to achieve mutual benefit by following it.

Notice that the Principle of Mutual Benefit never requires anyone to participate in a transaction from which she does not expect to gain. (It does not requires anyone to participate in *any* transaction.) However, an individual who has chosen to participate in a transaction is required to conform to the practice for that transaction, even if it would be in her self-interest to deviate. Thus, the principle requires adherence to practices of trust that are the normal expectation in that context – for example, the expectation that a buyer will pay the agreed price after the delivery of a good, or expectations about how a price will be adjusted when unforeseen contingencies occur.

If the Principle of Mutual Benefit is understood as the ethical underpinning for a network of opportunities for voluntary interactions, it has a natural counterpart in a principle that forbids each individual from obstructing other individuals' opportunities to cooperate with one another. For example, suppose that Annie wants to sell her car and that Bill and Charlie are potential buyers. If Bill makes a deal with Annie after having obstructed Charlie's efforts to make a bid, Annie's participation in that deal is not genuinely voluntary. An ethic of market behaviour that is based on the idea of cooperation for mutual benefit should proscribe anti-competitive actions that are intended to foreclose other agents' opportunities to transact with one another. Thus, Bruni and Sugden (2013: 156–157) include 'acceptance of competition' in a list of 'market virtues' derived from the principle that the *raison d'être* of the market is mutual benefit. Although many anti-competitive practices by firms – for example, collusion, dominance-creating mergers, investment strategies to disadvantage or exclude rivals – do not necessarily involve issues of transactional fairness between firms and consumers, some such practices do, because they work by imposing conditions on retail transactions which restrict consumers' opportunities to investigate the offers of rival firms.

3. The scope of transactional fairness

Transactional fairness is a property of individual transactions and not, as efficiency and income distribution are, properties of an economic system as a whole. A *transaction* is an economic interaction between specific *participants* (who may be individuals or firms), each of whom enters that interaction voluntarily. Considering any given transaction, the question that a concept of transactional fairness has to answer is: '*In this transaction, is each participant treating each other participant fairly?*' The formulation of this question immediately imposes restrictions on what can be relevant for assessments of transactional fairness. In this section, we formulate these restrictions as *scope requirements*. Although our main concern is with whether firms' pricing practices treat customers fairly, our discussion in this section will apply to transactional fairness in general. The concept of transactional fairness can apply to any participant in any transaction. (For example, one can ask whether a customer treats a firm fairly. Consider a consumer who buys an item of clothing on approval, wears it for one special occasion and then returns it as if unworn.)

We propose the following four scope requirements:

3.1. Irrelevance of externalities

Because transactional fairness is a property of the relationship between the participants to a transaction, external effects of that transaction on non-participants have no bearing on its fairness. Thus, in our example of bill shock, the benefit that Bella derives from Arthur's add-on bank charges is not relevant for an assessment of the fairness of the transaction between Arthur and his bank. Similarly, it might be a fact that, by increasing incentives to search, price walking tends to reduce the overall level of prices in the market for home insurance; but that fact would not be relevant to an assessment of the fairness of the transaction between Mr Smith and his insurance provider.

However, behaviour outside a particular transaction can be relevant for determining what kinds of behaviour count as fair in that transaction inasmuch as they affect normal expectations. For example, one might reasonably claim that the standards of fairness that apply to transactions between buyers and sellers at a car-boot sale or flea market are laxer than those that apply to transactions between a department store and its customers. Such a claim, as we interpret it, rests on the idea that standards of fairness are social norms, and that social norms can be context-dependent. Normal expectations about the behaviour of sellers in car-boot sales are descriptive of the behaviour of such sellers *in general*, but they affect what a buyer can expect in any *specific* car-boot transaction.

3.2. No requirement to incur losses

The ethical foundation of our concept of transactional fairness is a view of market transactions as cooperative. Thus, transactional fairness should not require individuals or firms to enter transactions from which they do not expect to gain relative to their outside options.

In this respect, transactional fairness is different from most of the concepts of pro-sociality that are represented in social preference theory. In the simplest such models, an individual with social preferences is altruistic – that is, willing to make trade-offs between personal benefits and benefits to other people (e.g. Becker, 1974). In more complex models, individuals are represented as having preferences for reductions in inequality between themselves and others (e.g. Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000), for increases in economic efficiency (e.g. Charness and Rabin, 2002), for confirmations of other people's expectations of benefit (e.g. Battigalli and Dufwenberg, 2007), or for the rewarding or punishing of other people for their kindness or unkindness, 'kindness' being interpreted as forgoing one's own benefits to benefit others (e.g. Rabin, 1993). In all these models, pro-

sociality is represented as a willingness to make some form of self-sacrifice. But self-sacrifice is inappropriate as a foundation for market ethics.

In the context of the market, agents treat one another fairly *within interactions that are directed at mutual benefit*; principles of transactional fairness impose constraints on the ways in which agents may properly seek to benefit themselves in such interactions. Thus, for example, bill-shock add-ons and loyalty penalties might be judged to breach those constraints, but the complaint against a firm that uses these pricing strategies is *not* that is failing to be altruistic towards its customers: fairness is expected, but altruism is not.

We recognise that there are exceptional situations in which a firm might be judged to have a moral obligation to engage in loss-making transactions. For example, one might think that an airline or train operating company ought to carry disabled passengers at its normal fares while providing them with the additional assistance they need, or that it ought to waive cost-recovering rebooking charges for customers who need to change their plans because of family emergencies. But we maintain that these are obligations of humanity or decency, not of transactional fairness. (We will say more about such obligations in Section 5.)

3.3 Non-paternalism

A welfare-based approach to normative economics that is addressed to a regulator or policy-maker might reasonably recommend paternalism in some situations – as advocates of libertarian and asymmetric paternalism indeed do. However, we maintain that a concept of fairness that is grounded on a view of market transactions as cooperative should treat individuals as responsible for their own choices.

A frequent claim in behavioural economics is that consumers can sometimes be made worse off by expansions of their choice sets. One version of this claim is the hypothesis of *choice overload*. Choice overload is said to occur when large choice sets induce low-quality decision-making by cognitively overloaded individuals, or reduce buyers' satisfaction with their final choices, or undermine individuals' motivation so that they avoid choosing altogether (e.g. Bown et al., 2003; Botti and Iyengar, 2006).[12] Choice overload is analogous to the external effects of transactions that we considered in Section 3.1. By adding a further product to a market in which many products are already on sale by other suppliers, a firm may increase the cognitive costs incurred by consumers in the market as a whole. But if the

[12] How far these hypotheses are supported by the balance of evidence is an open question: see the meta-analysis reported by Scheibehenne et al. (2010).

firm presents its own offer transparently, the costs that consumers incur in considering that offer are not unfairly imposed by the firm.[13]

A different version of the claim hypothesises *self-control failure* – that consumers are liable to choose superficially attractive products, contrary to their latent preferences or self-acknowledged long-term interests. In such cases, it is argued, consumers might be better off if certain options did not appear in their choice sets. This hypothesis is often represented in dual-self models in which a person's rational self (the 'Planner') can be subverted by an impulsive self (the 'Doer'). Thaler and Sunstein (2008: 41–49) endorse this model and give the example of a Cinnabon stand at Chicago O'Hare Airport, whose oven aromas subvert the health-oriented intentions of Planners heading for the nearby fruit and yoghurt stand. Given evidence of self-control failure, a sufficiently paternalistic social planner might invoke welfare-based justifications for imposing restrictions on the offers that firms make to consumers. But would such restrictions be justified *on grounds of transactional fairness*? We think not. We take it that the Cinnabon stand is offering consumers an opportunity to buy a familiar product at a familiar price. The oven aromas are not deceptive; they remind potential customers of the actual qualities of the experience of consuming the product. We maintain that this should be understood as a fair offer. The customer's decision to buy a cinnamon bun may be impulsive, but it is *her* decision. It is fair that she takes responsibility for it.

3.4 Availability of information about co-participants

Transactional fairness is concerned with ethical standards that can be action-guiding for individual firms and consumers. Principles of transactional fairness must therefore be formulated in terms of information that is available to the firms and consumers who are expected to be guided by them.

Notice the contrast with efficiency- and welfare-based approaches to normative economics, which are intended to guide the actions of regulators or policy-makers. Because efficiency and social welfare are properties of a whole economy, and because an economy is a complex system, individual firms and consumers cannot be expected to base their decisions directly on such criteria. Instead, the expectation is that normative economics provides

[13] Different issues are involved if a firm puts multiple offers on the market, with the intention that some consumers will fail to notice the offers that are best for them. We discuss such cases in Section 4.2.

guidance about how to design institutional mechanisms, or how to set parameters such as tax rates, so that efficiency or welfare emerges as an unintended outcome of decisions by firms and consumers who act on more immediate criteria. In contrast, transactional fairness must restrict itself to outcomes that are intended by – or at least, reasonably foreseeable by – decision makers.

Our feasibility requirement is that each participant in a transaction should be able to assess the fairness of his (or its) own behaviour in that transaction. Although the external enforcement of fairness principles – by social pressure, naming and shaming in the media, or by regulation – requires that violations can sometimes be observed by others, it is not essential that each participant in a transaction always has sufficient information to allow him to assess whether the behaviour *of other participants* is fair. Indeed, many practices that are commonly judged unfair are based on information asymmetries between participants in a transaction. For example, think of an insurance firm that quotes high renewal prices to long-standing customers without telling them about their eligibility for its own cheaper tariffs. Or think of the consumer who returns 'as new' the dress she has bought on approval and worn at a party.

The condition that agents are able to assess the fairness of their own actions imposes limits on how far transactional fairness can require firms to meet their customers' preferences. Under the assumption that consumers have integrated preferences that are revealed in their decisions, a firm might be expected to know the main population-level properties of consumer preferences with respect to its own products. In some contexts, more might be expected. For example, professional codes may require that clients' preferences are solicited as part of a transaction (e.g. financial advisers asking about risk preferences). But in many retail settings the firm cannot know the actual preferences of particular customers. Thus, while it would be unfair for a firm to use the small print of an offer to hide information about an add-on that almost all customers could be expected to incur (e.g. a standard delivery charge), it is not possible to describe *every* add-on with a degree of salience that matches its particular importance to *every* customer. The information problem becomes more severe if there are naïve consumers whose behaviour does not reveal integrated preferences. If latent preference is not an empirical concept – a possibility that we considered in Section 1 – transactional fairness clearly cannot require firms to take account of customers' latent preferences.

What *can* a firm be expected to know about its customers? Some basic properties of preferences, such as that most consumers prefer to pay less rather than more for given products, or that buyers of electrical appliances prefer these to work when plugged in, are uncontroversial. Less obviously, the fact that a consumer has *chosen* to participate in a particular transaction with a firm has significant information content. The firm knows that the consumer made that decision in the light of the information at her disposal. Some of the latter information was supplied by firm itself, and so is known to the firm. In addition, the firm might reasonably assume that its customers have some knowledge about normal practices in the market in which they are participating. For example, introductory offers may be common practice in a market with consumer switching costs or for experience goods and services. Similarly, a seller in a car-boot sale can reasonably assume that he is selling to someone who knows that trades in car-boot sales are made on *caveat emptor* terms.

4. Transactional fairness in firms' pricing practices

Our aim is to propose a definition of transactional fairness that can be used in assessing the fairness of firms' pricing practices.[14] This conception is grounded in the ethic of mutual benefit, as presented in Section 2, and is compatible with the scope requirements set out in Section 3. We begin with a summary definition:

> *Transactional Fairness.* Transactional fairness requires that a firm acts in such a way that consumers with normal expectations about pricing practices in the relevant market are able to understand the consequences of transacting with that firm (*No Deception*) and are not hindered from terminating a relationship with the firm or from transacting with alternative sellers (*No Hindrance*). It also requires that the firm is able to explain the rationale of its pricing practices, locating them as part of a business model based on mutual benefit between the firm and its customers, and is willing to provide the explanation publicly (*Public Explanation*).

We now explain the components of this definition in more detail.

---

[14] In relation to any given transaction, we treat 'the firm' as the trading entity that is recognised by consumers. For example, NatWest, Royal Bank of Scotland, Adam & Co. and Coutts are brands that are all part of NatWest Group. Consumers' normal expectations on pricing and service are likely to be associated with the individual brands rather than the group. For other issues, such as asset protection, the group would be relevant.

4.1 Normal expectations

'Normal expectations' were defined in Section 2. Normal expectations about pricing practices can be interpreted as default settings for the terms of implicit contracts between firms and consumers. (Those settings can be over-ridden, but only with the explicit agreement of both parties.) There is an implicit contract between a firm and a customer if they are mutually aware of (i) the market context for the transaction they have initiated, (ii) the business (and consumer) practices that normally apply in this context and that are obviously relevant to the participation decision, and (iii) the fact that participation in the transaction has been freely chosen. According to the Principle of Mutual Benefit, the parties to an implicit contract should comply with its terms, even if advantage could be gained by deviating. Our definition of transactional fairness incorporates that requirement.

Implicit contracts, based on (often market-specific) expectations as to what are acceptable commercial practices, are common in business-to-business transactions. For example, there may be hard bargaining after which an agreement may be secured by a handshake. A detailed contract usually follows for high-value, one-off transactions, but relational contracting, with important unwritten conditions, often suffices in low-value or ongoing relationships.[15] These norms reduce transaction costs, facilitate flexible adaptation to changing circumstances, support rapid trading and allow parties to focus their detailed attention on other matters such as their core production. They can be enforced by reputation, withdrawal of future business and resort to contract law.[16] The scale and value of business transactions make these mechanisms credible.

Implicit contracts between firms and consumers are similar in principle but often less enforceable. (If each consumer accounts for only a small share of the firm's sales, the value of an ongoing relationship with any individual customer is correspondingly small.) Nevertheless, contravening an implicit contract is unfair. If expectations can be expected to

---

[15] See Macneil (1978). For a summary of the evidence, see the introduction to Baker et al (2002).

[16] The legal standing of relational contracts is mixed at best, and tied to the concepts of good faith and fair dealing. Most civil law systems have an overriding principle that contracting parties should act in good faith, and the concept can also be found in commercial codes in many common law systems. (The UK is generally considered an exception, although good faith is mentioned in the Consumer Rights Act 2015 s.62(4) in relation to 'a significant imbalance in the parties' rights and obligations'.) However, 'the notion of good faith… actually means different things both *within* a particular legal system and *between* the legal systems' (Whittaker and Zimmermann, 2000, p.690). The result is that little can be said about what good faith actually means 'and what can be said is not very helpful for deciding concrete cases' (ibid p.30).

differ between consumers (for example, at times when new business models are emerging, or when the context of a transaction can be interpreted in different ways, implying different normal expectations), a firm should make public any business practices that it intends to follow and that are obviously relevant to the participation decision. If a firm relies on (and does not try to correct) a customer's mistaken beliefs about its business practices, it is engaging in passive deception.

In interpreting 'normal expectations', some allowance must be made for vulnerable consumers. We will say that (in relation to a specific transaction) a consumer is *transactionally vulnerable* if her capacity to make considered and well-informed decisions is impaired by factors outside her or the firm's control. 'Impairment' may be due to long-lasting cognitive limitations (e.g. young children, individuals with dementia) or temporary distress (e.g. a recently bereaved person planning a funeral). Vulnerability with respect to information could also be due to inability to access commonly-used sources of information (e.g. lack of internet access). Insofar as the firm can recognise transactionally vulnerable customers, it should interpret its implicit contracts with them in accordance with what those customers can be expected to know and understand. Because the requirements that transactional fairness imposes on a firm are conditional on the firm's information, vulnerability is relevant for transactional fairness only to the extent that the firm can recognise it. There are very substantial privacy issues in identifying vulnerable individuals. Nevertheless, some natural correlates of vulnerability are easily identified (e.g. age), and some sales situations may naturally present other evidence (e.g. severe mental impairment is likely to be revealed in person-to-person sales, bereavement is revealed in funeral planning).

4.2. No Deception

There is deception, and therefore transactional unfairness, if a consumer is enticed into a transaction by a firm's use of misleading information, or by its hiding obviously relevant information about the transaction. 'Hiding' includes presenting information with misleading salience that draws attention away from what is likely to be most relevant. This much is the widely accepted and familiar territory of robust consumer law.

'Obviously relevant information' includes information about the prices of add-ons that are effectively unavoidable (e.g. delivery charges, booking fees, fees for debit card payments for online sales) or that most customers would expect to buy in combination with the main product (e.g. product-specific ink cartridges for printers, product-specific brush

heads for electric toothbrushes). Where add-ons are relevant only for relatively small numbers of customers (e.g. additional charges for delivery to particularly remote locations) or would be needed only in unlikely contingencies (e.g. fees for replacing lost travel documents), it is not unfair for their prices to be relegated to the small print of an offer, but for this kind of low-visibility price to be non-deceptive, it should be in line with normal pricing practices in the relevant market.

It is unfair for a firm to hide information about other relevant tariffs that it offers and for which the consumer is eligible: this information is 'obviously relevant' to the transaction between the specific firm and the specific consumer. This rules out forms of price discrimination that rely on consumers' lack of information about *the firm's own* prices.[17] For example, it would be unfair for a train operating company to sell 'any time' tickets to walk-up customers in off-peak periods without informing them about its cheaper off-peak tickets. If, instead of quoting take-it-or-leave it prices, the firm is stating an initial offer that it is willing to negotiate, or that it is willing to reconsider if a potential customer can show that a rival has quoted a lower price, it should make this clear as a property of that offer. Not doing this can sometimes be a profitable form of price discrimination, based on differences in consumers' knowledge about the firm's pricing strategy; but it contravenes the No Deception condition.

Moving into the domain of competition law, a firm is under no legal obligation to tell a potential customer about *its competitors'* offers. One might ask whether it is fair (and not merely legal) for a firm to sell at prices that it knows are higher than those of a rival, and without announcing that fact.[18] For example, GAP (guaranteed asset protection) insurance is often sold as an optional add-on when a new car is bought; typically, an equivalent product could be bought as a stand-alone purchase at a significantly lower price (FCA, 2018a).[19]

---

[17] Recall from note 9 that 'the firm' may be identified by its trading name rather than its ownership. Thus, our proposal does not prohibit forms of price discrimination in which what is effectively the same product is sold at different prices under different brand names, without the consumer being informed of this fact.

[18] A related issue was debated by ancient and medieval philosophers. A merchant is carrying wheat to a city where grain is in short supply and the price is high. He knows that other sellers of wheat will arrive soon, and so the price will fall. His potential customers do not know this. Does justice require the merchant to reveal this information? Aquinas (1265–74, Part II.II, Question 77, Article 3) argues that it does not.

[19] The evidence suggests that, for many consumers, willingness to pay for GAP insurance is context-dependent: a person who chooses not to buy it as an add-on in the saleroom is unlikely to buy a much cheaper stand-alone product after driving the car home (FCA, 2018a). Is this the result of

Clearly, the car saleroom has a point-of-sale advantage in the insurance market, which might reasonably be judged anti-competitive and which a regulator might try to remedy.[20] Nevertheless, we maintain that not revealing a rival's price is compatible with transactional fairness. From the customer's point of view, information about the rival firm's price is undoubtedly relevant for her choice between them; but it is not information *about the transaction that is being judged fair or unfair*. Taking a wider view, 'each firm reveals its competitors' prices' could not persist as a normal expectation in a competitive market. Such a practice would undermine both consumer responsibility to search and rivalry between firms – mechanisms that are fundamental to the working of the market. (Recall the discussion of price dispersion in Section 1.)

The position we defended in the preceding paragraph does not necessarily apply when consumers are transactionally vulnerable. Take the case of an unscrupulous firm that specialises in doorstep contact with elderly householders, offering to do repair work at prices far above market levels. If these customers are known to lack the mental resources necessary for effective price search, this practice should be viewed as transactionally unfair.[21]

As we noted in the introduction, there is an increasing tendency for products to be sold on indefinite or default contracts (e.g. banking, domestic energy, mobile phones, wifi, home entertainment), for rental deals to replace sales of durable goods (e.g. cars, computer software), and for subscription contracts to replace payment-per-item selling (e.g. sales through Amazon Prime, streaming services).[22] Such business models involve an ongoing relationship between the firm and 'its' customer (or 'client'). This creates a continuing implicit contract which expands the scope of the No Deception condition.

---

psychological bias in the saleroom (over-attention to the enticing features of the new car, susceptibility to saleroom pressure) or psychological bias afterwards (lack of engagement with financial matters, procrastination)? Compare our discussion of 'bias' in Section 1.

[20] In fact, the relevant regulator did take action (FCA, 2018a).

[21] This issue has a philosophical pedigree too. In an early discussion of price discrimination, Kant (1785/ 2002: 13) claims that it is 'in conformity with duty that the merchant should not overcharge his inexperienced customers', and gives the example of a transaction with a child. (Kant notes that duty and self-interest would coincide for a merchant in a sufficiently competitive market, but even if the merchant has monopoly power, the duty of non-discrimination remains.)

[22] To some extent, this shift is a return to business models that were common before the supermarket era. It was once common for households to have ongoing relationships with specific suppliers of milk, groceries, fish and meat products, sometimes with home deliveries and purchase on credit.

The No Deception condition implies that there is transactional unfairness if a firm attempts to retain an existing consumer by giving misleading information, or by hiding obviously relevant information about the terms of the continuing transaction. Thus, it is unfair if the firm does not periodically provide an existing customer with relevant information about changes in its prices, or about the absence of a price change for a product whose cost of supply is falling over time. Several UK regulators have recently adopted policies requiring firms to provide such information when contracts are renewed (e.g. the requirement that insurance renewal documents report the customer's previous price alongside the renewal quote).

As in the case of pre-purchase fairness, it is unfair for a firm to hide information about alternative tariffs that it offers and for which the customer is eligible, or about its willingness to renegotiate initial renewal quotes. These principles rule out many forms of price discrimination between 'front book' (recently acquired) and 'back book' (long-standing) customers: fairness requires the firm to inform its back-book customers about their eligibility for front-book offers. It is not necessarily unfair for a firm to make low-price introductory offers that are available only to new customers. Such offers are compatible with intentions for mutual benefit if the low price is paid only for an introductory period, and if the intention is to allow new customers to sample the firm's product, or to compensate them for costs of searching and switching. However, it is contrary to the No Deception condition to mislead consumers about the ease with which they will be able to cancel at the end of the introductory period. (As we explain in Section 4.3, such a practice would also fall foul of the No Hindrance condition.)

4.3 No Hindrance

The No Hindrance condition requires that consumers with normal expectations are not hindered from terminating a relationship with a firm or from transacting with another firm. The first clause is in the spirit of consumer law; the second is in the spirit of competition law.

There is transactional unfairness if a firm uses pricing practices that deter consumers from searching for competitors' offers. For example, industrial economists and regulators have analysed firms' use of *time-limited* (or *exploding*) offers – offers that must be accepted or rejected within a time frame that is too short to allow a potential buyer to search for other offers. A related (alleged) practice in internet selling is to use cookies that record individual consumers' search behaviour and to quote higher prices to consumers who return to a firm's

website after searching elsewhere. By creating barriers to search, such practices tend to raise prices (e.g. Office of Fair Trading, 2010; Armstrong and Zhou, 2016). Irrespective of that tendency, pricing practices that systematically penalise consumers for investigating other firms' offers are transactionally unfair. Selling strategies that make posted prices and available offers liable to change at short notice (e.g. because capacity constraints are reached, or a fixed stock of a good is sold out) are not inherently unfair. Correspondingly, it is not inherently unfair if a consumer misses out on an offer by searching for, and failing to find, something better. But it is unfair for a firm to make an offer that is *conditional on* the consumer's not making further searches.

In the context of ongoing relationships between firms and customers, there is transactional unfairness if a firm attempts to retain an existing customer by making it difficult for her to cancel a contract for which renewal is the default. Familiar hindrances to cancellation include procedures that require consumers who signed up online to cancel by mail or phone (sometimes using phone lines with very slow answering services), and online interfaces in which cancellation options are not easily visible or involve unnecessary sequences of operations. A good rule of thumb is *exit/entry equivalence* – that 'consumers should find it as easy to exit a contract as it was to enter' (CMA, 2019: para. 130). This should be read as referring to indefinite contracts with continuous or recurring opportunities to choose between renewal and cancellation; it does not release consumers from fixed-term contracts that were fairly entered into (e.g. fixed-term mortgages with pre-specified interest rates).

Analogously with passive deception, there can be unfair 'passive hindrance'. For example, if a contract is subject to periodic renewal by direct debit, it is unreasonable to expect consumers to remember renewal dates in the absence of reminders or renewal statements. Fairness requires a firm to give customers sufficient notification of upcoming renewal dates to allow them to search for alternative offers.

4.4 Public Explanation

The Public Explanation condition requires that a firm is able to explain the rationale of its pricing practices and is willing to provide this explanation publicly. By the 'rationale' of a pricing practice, we mean *the firm's reasons* for choosing to use it. There are two ethical constraints here. First, the reasons must be genuine explanations of the firm's behaviour. Second, the rationale must be in terms of a business model based on mutual benefit between

the firm and its customers. By 'mutual benefit', we mean that the firm pursues the interests of its owners within the constraints of No Deception and No Hindrance. In other words, the firm should be able to explain the practice as involving neither deception nor hindrance. Subject to those constraints and in relation to an assessment of transactional fairness, profit-maximisation is a legitimate rationale of a pricing practice. There is no requirement of paternalism on the part of the firm: a firm treats its customers as benefitting from their transactions with it if they enter and leave those transactions voluntarily, without being subject to deception or hindrance. The added value of the Public Explanation condition is that it requires firms to declare the intentions behind their actions. Intentions can be important when a pricing practice has more than one possible explanation.

For example, consider the phenomenon of 'paying not to go to the gym'. Della Vigna and Malmendier (2004, 2006) investigate consumers' buying behaviour when gyms offer a choice between pay-as-you-go and membership tariffs. They find that, on average, people who choose monthly membership tariffs pay around 70 per cent more than the total pay-as-you-go prices of the visits they actually make. Della Vigna and Malmendier (2006: 716) conclude that the best explanation for this is that gym users tend to over-estimate how many gym visits they will make in the future.[23] In a competitive market, consumers whose forecasts are correct (savvies) are cross-subsidised by those whose forecasts are not (naïves). A similar analysis applies to low-price mobile phone tariffs with low usage thresholds and high 'overage' charges: in this case, savvies are cross-subsidised by naïve *under*-forecasters (Grubb, 2015).

In these examples, there are (at least) two possible explanations for the existence of the tariffs that are chosen by mis-forecasting naïves. One explanation is that the firm is using pricing strategies that are designed to discriminate between consumers with different *actual* usages, on the assumption that each consumer will choose the tariff that is cheapest for her. Under this assumption, the firm offers multiple tariffs with the aim of maximising profit, but with no intention to deceive. It is unavoidable that some consumers who mis-forecast their usage end up paying more than they needed to have done, but that is an unintended by-product of a pricing strategy in which each tariff is offered as a good buy for its intended purchasers. An alternative explanation is that the tariffs chosen by mis-forecasting naïves are

[23] The evidence supports this explanation rather than one that is often suggested by behavioural economists – that membership is a self-control device by which consumers pay in advance for visits, irrespective of whether they will in fact be made.

put on the market with the primary intention that they will be chosen in exactly this way. Thus, for example, the apparently low price of the low-usage phone tariff might be merely a bait to attract overage payments. In that case, the firm's pricing strategy would be based on passive deception.

By requiring a firm to be willing to explain the rationale of its pricing strategies, the Public Explanation condition puts some pressure on the firm to avoid transactional unfairness and facilitates regulation. We will say more about this in Section 6.

5. Discrimination between consumers

In the preceding section, we categorised various forms of price discrimination as transactionally unfair. Many of those pricing practices (e.g. hidden add-ons, price walking, overage charges designed to trap under-forecasting consumers, time-limited offers, barriers to the cancellation of contracts, not informing customers about the firm's own tariffs) discriminate against naïve consumers and in favour of those who are savvy. In our analysis, however, the unfairness of such practices derives from properties of deception or hindrance that are located in the transaction between the firm and the naïve consumer, and not in the difference between the firm's treatment of the two classes of consumers. This is fundamental to our concept of transactional fairness: we have defined transactional fairness as a property of individual transactions.

To repeat what we said in the introduction, we are not proposing transactional fairness as the *only* normative criterion for assessing pricing practices: we see it as complementary with the standard criteria of efficiency and distributional equality. In the 'transactional' sense of fairness, price discrimination is not intrinsically unfair. In terms of standard criteria, price discrimination by profit-seeking firms can be pro-competitive (e.g. when a firm uses it to enter a market in which a rival has a dominant position) or anti-competitive (e.g. when a dominant firm uses it to exclude potential entrants). It can be distributionally progressive (e.g. when the rich have less elastic demand than the poor) or regressive (e.g. when richer consumers can access a greater number of potential suppliers, and so have more elastic demand). When used by a monopolist, price discrimination can increase economic efficiency by enabling a product to be sold to consumers whose willingness to pay is relatively low, and it can reduce consumer welfare by extracting more surplus than is needed to recover fixed costs.

When price discrimination is used to recover fixed costs, there are significant parallels between normative criteria for price discrimination and principles of just taxation. In the traditional literature of public finance, two principles of just taxation are often discussed – the *ability-to-pay principle* and the *benefit principle* (e.g. Musgrave, 1959). Applied to price discrimination, the ability-to-pay principle implies that richer consumers of a given product should pay more than poorer consumers, as is approximated when concessionary prices are offered to buyers who are in full-time education, retired or 'unwaged'. The benefit principle implies that consumers who benefit more from a given product should pay more than those who benefit less. This fits naturally with the idea of market transactions as cooperation for mutual benefit (Sugden, 2018: 160–164). However, we do not claim that any particular relationship between price and benefit is a requirement of transactional fairness.

In many cases, profit-seeking price discrimination by firms is broadly in line with the benefit principle. For example, consider how the prices charged by (non-budget) airlines for round-trip tickets differ according to the interval between outward and return flights. (Significantly, this discrimination does not reflect cost differences: it not based on the separate flights that the consumer buys, but on which flight is bought in combination with which.) The rationale is that business travellers have low price elasticity of demand and are particularly likely to choose short stays. The result is that consumers with higher willingness to pay bear a higher share of the airline's fixed costs. A similar analysis applies to introductory offers for experience goods, as discussed in Section 4.2: consumers who are not familiar with a firm's (assumedly good quality) product have lower willingness to pay for it than do the firm's current customers.

We further recognise that price discrimination can be unfair in ways that are not 'transactional', and that the standard economic criteria of efficiency and distribution do not take into account. That people should be treated equally, irrespective of 'protected' characteristics such as age, disability, sex, ethnicity, religion and sexual orientation, is a very widely held ethical principle, upheld by law in most democracies.[24] If the *raison d'être* of the market is mutual benefit, a disposition to make mutually beneficial transactions with others on terms of equality, whoever those others may be, is a 'market virtue' (Bruni and Sugden, 2013: 154). Nevertheless, pricing practices that discriminate on the basis of

[24] For example, the UK's Equality Act 2010 consolidated earlier anti-discrimination law by setting out a list of 'protected characteristics': age, disability, gender reassignment, marriage and civil partnership, pregnancy and maternity, race, religion or belief, sex, and sexual orientation. The Act prohibits direct and indirect discrimination based on these protected characteristics.

protected characteristics need not be *transactionally* unfair. Problem cases arise when protected characteristics are correlated with factors that would otherwise be legitimate bases for price discrimination. For example, concessionary prices for the young and for the old are based on correlations between age and income; differential insurance premia for male and female drivers (common until a European Court of Justice ruling in 2012) were based on correlations between gender and accident risk. In not addressing these difficult issues, we should not be thought to be denying their importance. We believe that clarity is best served by treating transactional fairness as a distinct form of fairness.

6. The role of regulation

The underlying ethical idea of 'mutual benefit', as presented in Section 2, is that a well-functioning market is a network of cooperative interactions that individuals enter voluntarily, knowing what to expect, and remain free to leave. This idea suggests two guiding principles.

First, regulation should facilitate the formation and maintenance of trading practices that allow consumers to predict, as precisely as possible, the consequences of entering a transaction. Second, regulation should not hinder the emergence of (non-deceptive and non-hindering) business models which, by offering new opportunities to consumers, may create new expectations. For example, until relatively recently, airlines' posted prices varied according to stated conditions (e.g. the day of the week on which flights were made) but price schedules were changed only infrequently. Digital technology now allows airlines to increase load factors by using 'dynamic pricing', adjusting prices minute by minute to reflect current demand for specific flights. When these pricing strategies were first introduced, they were contrary to consumers' normal expectations, but a regulation that prohibited short-term changes in price schedules would have cut off opportunities for mutually beneficial transactions. Notice that, in line with the ethic of mutual benefit, these guiding principles do not assume that consumers have integrated preferences (whether revealed or latent) and do not require regulators to make paternalistic judgements about consumers' best interests.

If a firm's pricing practices are transactionally unfair, they have harmful effects on the customers who are treated unfairly, and can persist only as long as those consumers continue to patronise the firm. Thus, in principle, firms might find it profitable to compete in developing reputations for transactional fairness. There are many historical examples of consumer-facing firms that have built and maintained market share by this reputational route. However, this is likely to be an imperfect mechanism, particularly at times when the retail

sector is experiencing rapid change, making it harder for consumers to recognise and keep track of firms' behaviour. As we noted in the introduction, there can be externalities in reputation: after experiencing instances of unfairness, consumers may draw inferences about the unfairness of firms in general, and not merely about the unfairness of specific firms.[25] This suggests a role for regulation in defining and maintaining standards of transactional fairness.

Principles of non-deception and non-hindrance are already written into consumer and competition law, but generally with the emphasis on active deception and active hindrance. What we have called passive deception and passive hindrance rest in a grey area of business practices that are unfair without being clearly illegal, except insofar as they contravene specific regulations. Such regulations are often designed as remedies for existing practices that have been judged to be unfair. There is an entrepreneurial dynamic here, with an inbuilt time lag. Firms discover new strategies (perhaps made possible by advances in technology) for making profit within the constraints imposed by existing regulations; when those strategies are judged unfair, new regulations are introduced.

Our aim in this paper has been to contribute to the development of general principles that can guide regulators in assessing the fairness or unfairness of firms' pricing practices. If (as we have proposed) those principles use only information that is directly available to firms, firms should be able to predict whether or not particular new practices would be permitted by the regulator. This facilitates both compliance and, where necessary, enforcement.

One regulatory role might be to prohibit specific types of passive deception in a category of markets (e.g. a firm 'hiding' prices from an eligible consumer, or 'price walking' for a consumer on auto-renewal). Similarly, certain types of passive hindrance may be prohibited in relevant markets (e.g. making it harder to leave a contract than to sign-up, or 'rolling-over' fixed-term contracts without sufficient warning). UK regulators have been

[25] Tirole (1996) develops a theory of reputational externalities to show that 'after episodes of bad behaviour, either the group is stuck in a bad-reputation steady state, or trust takes several periods to re-establish' (p.18). Furthermore, 'an increase in product market competition may make it difficult for firms to sustain their reputation' (p.3). A recent empirical test investigates the firm-specific scandal of VW's deliberately misleading behaviour in its conduct of diesel emissions tests, and the negative externality this had on other German car manufacturers in the US market. Bachmann et al. (2017) find that German firms unconnected to VW Group (e.g. BMW, Daimler) suffered substantial loss of market value and sales, including for their petrol vehicles. They also suffered a deterioration of sentiment (a measure of approval/reputation) on Twitter.

active in trying to address such issues, but consistency is difficult in the absence of guiding normative economic principles.

A second role would be to administer the Public Explanation condition. This requires firms to cooperate with regulators in the public process of ensuring that business practices are fair. Regulatory agencies can be a repository for complaints of unfairness made by consumers and advocacy groups. Regulators can ask firms to explain their practices, ask for verification of questionable claims, make public comment on their findings, and potentially enforce a change in business practice.

Because the concepts of deception and hindrance are defined in relation to normal expectations, there is a third role for regulation in shaping those expectations. For example, we have argued that 'hidden' add-ons are unfair if consumers would normally expect them to be included in the headline price. But in a market in which different firms define their headline prices in different ways, there may be no precise 'normal expectation', with the implication that no individual firm's practice is deceptive. A regulation that imposes a specific definition (e.g. requiring that headline prices are stated inclusive of taxes and standard delivery charges) can *create* a normal expectation and a corresponding category of transactional unfairness.

It is reasonable to hope that the regulation needed to secure transactional fairness would not be onerous. There is currently a lack of corporate understanding about what is acceptable practice in pricing. We suggest that clear, principled guidance on transactional fairness would make it easier for firms to act fairly.

References

Aquinas, Thomas (1265–1274). *Summa Theologica*. Available at: http://www.documentacatholicaomnia.eu/03d/1225-1274,_Thomas_Aquinas,_Summa_Theologiae_%5B1%5D,_EN.pdf , accessed 4 August 2020.

Ariely, Dan, George Loewenstein and Drazen Prelec (2003). Coherent arbitrariness: stable demand curves without stable preferences. *Quarterly Journal of Economics* 118: 73–105.

Armstrong, Mark (2015). Search and ripoff externalities. *Review of Industrial Organization* 47: 273–302.

Armstrong, Mark and John Vickers (2012). Consumer protection and contingent charges. *Journal of Economic Literature* 50(2): 477–493.

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

Armstrong, Mark and Jidong Zhou (2016). Search deterrence. *Review of Economic Studies* 83: 26–57.

Bachmann, Rüdiger, Gabriel Ehrlich and Dimitrije Ruzic (2017). Firms and collective reputation: The Volkswagen emissions scandal as a case study. *CEPR Discussion Paper* No. 12504.

Battigalli, Pierpaolo and Martin Dufwenberg (2007). Guilt in games. *American Economic Review: Papers and Proceedings* 97: 171–176.

Becker, Gary (1974). A theory of social interactions. *Journal of Political Economy* 82: 1063–1093.

Bernheim, Douglas (2016). The good, the bad, and the ugly: a unified approach to behavioural welfare economics. *Journal of Benefit-Cost Analysis*. Available on CJO 2016 doi:10.1017/bca.2016.5.

Bershears, John, James Choi, David Laibson and Brigitte Madrian (2008). How are preferences revealed? *Journal of Public Economics* 92: 1787–1794.

Bicchieri, Cristina (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms.* Cambridge: Cambridge University Press.

Bolton, Gary and Axel Ockenfels (2000). ERC: A theory of equity, reciprocity and competition. *American Economic Review* 90: 166–193

Bordalo, Pedro, Nicola Gennaioli and Andrei Shleifer (2013). Salience and consumer choice. *Journal of Political Economy* 121: 803–843.

Botti, Simona and Sheena Iyengar (2006). The dark side of choice: when choice impairs social welfare. *Journal of Public Policy and Marketing*, 25: 24–38.

Bown, Nicola, Daniel Read and Barbara Summers (2003). The lure of choice. *Journal of Behavioral Decision Making* 16: 297–308.

Bruni, Luigino and Robert Sugden (2013). Reclaiming virtue ethics for economics. *Journal of Economic Perspectives* 27: 141–164.

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.

Charness, Gary and Matthew Rabin (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117: 817–869.

Cialdini, Robert, Raymond Reno and Carl Kallgren (1990). A focus theory of normative conduct: recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology* 58: 1015–1026.

Competition and Markets Authority [CMA] (2016a). *Energy market investigation: final report*. 24 June 2016.

Competition and Markets Authority [CMA] (2016b). *Retail banking market investigation: final report*. 9 August 2016.

Competition and Markets Authority [CMA] (2018). *Tackling the loyalty penalty: response to a super-complaint made by Citizens Advice on 28 September 2018*. 19 December 2018.

Competition and Markets Authority [CMA] (2019). *Consumer Vulnerability: Challenges and Potential Solutions*. February 2019.

Della Vigna, Stefano and Ulrike Malmendier (2004). Contract design and self-control. *Quarterly Journal of Economics* 119: 353–402.

Della Vigna, Stefano and Ulrike Malmendier (2006). Paying not to go to the gym. *American Economic Review* 96: 694–719.

Diamond, Peter (1971). A model of price adjustment. *Journal of Economic Theory* 3(2): 156–168.

Dupuit, Jules (1844/ 1952). On the measurement of the utility of public works. *International Economic Papers* 2 (1952): 83–110. Translated by R.H. Barback. First published in French in 1844. (Page references to reprint in *Cost–Benefit Analysis*, edited by Richard Layard. Penguin, 1973.)

European Union (2005). *Unfair Commercial Practices Directive*. Directive 2005/29/EC.

Fehr, Ernst and Klaus Schmidt (1999). A theory of fairness, competition and cooperation. *Quarterly Journal of Economics* 114: 817–868.

Fehr, Ernst and Simon Gächter (2000). Cooperation and punishment in public goods experiments. *American Economic Review* 90: 980–94.

Financial Conduct Authority [FCA] (2015) *Cash savings market study report: final findings*. Market Study 14/2.3, January 2015.

Financial Conduct Authority [FCA] (2018a). *An evaluation of our guaranteed asset protection insurance intervention*. Evaluation Paper 18/1.

Financial Conduct Authority [FCA] (2018b) *Fair Pricing in Financial Services*. Discussion Paper 18/9.

Financial Conduct Authority [FCA] (2019) *General insurance pricing practices: Market Study interim report* MS18/1.2.

Gabaix, Xavier and David Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121(2): 505–540.

Grubb, Michael (2015). Consumer inattention and bill-shock regulation. *Review of Economic Studies* 82(1): 219–257.

Infante, Gerardo, Guilhem Lecouteux and Robert Sugden (2016). Preference purification and the inner rational agent: a critique of the conventional wisdom of behavioural welfare economics. *Journal of Economic Methodology* 23: 1–25.

Isoni, Andrea, Robert Sugden and Jiwei Zheng (2020). Voluntary Interaction and the Principle of Mutual Benefit. [currently a working paper, revised and resubmitted to *Journal of Political Economy*.]

Kant, Immanuel (1785/ 2002). *Groundwork for the Metaphysics of Morals*. Translated by Allen Wood. New Haven, CT: Yale University Press.

Kőszegi, Botond and Matthew Rabin (2007). Mistakes in choice-based welfare analysis. *American Economic Review* 97: 477–481.

McQuillin, Ben and Robert Sugden (2012). How the market responds to dynamically inconsistent preferences. *Social Choice and Welfare* 38: 617–634.

Musgrave, Richard (1959). *The Theory of Public Finance.* New York: McGraw-Hill.

Office of Fair Trading (2010). *Advertising of Prices*. Report OFT 1291. London: Office of Fair Trading.

Rabin, Matthew (1993). Incorporating fairness into game theory and economics. *American Economic Review* 83: 1281–1302.

Rizzo, Mario and Glen Whitman (2020). *Escaping Paternalism: Rationality, Behavioural Economics and Public Policy*. Cambridge: Cambridge University Press.

Salop, Steven and Joseph Stiglitz (1977). Bargains and ripoffs: a model of monopolistically competitive price dispersion. *Review of Economic Studies* 44(3): 493–510.

Scheibehenne, Benjamin, Rainer Griefeneder and Peter M. Todd (2010). Can there ever be too many options? A meta-analytic review of choice overload. *Journal of Consumer Research* 37: 409–425.

Stango, Victor and Jonathan Zinmann (2014). Limited and varying consumer attention: evidence from shocks to the salience of bank overdraft fees. *Review of Financial Studies* 27: 990–1030.

Stiglitz, Joseph (1979). Equilibrium in product markets with imperfect information. *American Economic Review* 69(2): 339–345.

Sugden, Robert (2004). The opportunity criterion: consumer sovereignty without the assumption of coherent preferences. *American Economic Review* 94: 1014–1033.

Sugden, Robert (2015). Looking for a psychology for the inner rational agent. *Social Theory and Practice* 41: 579–598.

Sugden, Robert (2018). *The Community of Advantage: A Behavioural Economist's Defence of the Market*. Oxford: Oxford University Press.

Sunstein, Cass and Richard Thaler (2003). Libertarian paternalism is not an oxymoron. *University of Chicago Law Review*, 70: 1159–1202.

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

Varian, Hal (1980). A model of sales. *American Economic Review* 70(4): 651–659.