

The economics of online personalised pricing

May 2013

OFT1488

© Crown copyright 2013

This report has been written by Patrick Coen and Natalie Timan of the OFT, and has benefitted from comments by various other members of staff at the OFT, including Amelia Fletcher, Kirsten Edwards-Warren, Hugh Mullan and Leanne Graveney. The OFT would also like to thank Steffen Huck (University College London and Wissenschaftszentrum Berlin), Alexandre de Cornière (University of Oxford) and various seminar participants for their helpful comments.

You may reuse this information (not including logos) free of charge in any format or medium, under the terms of the Open Government Licence. To view this licence, visit www.nationalarchives.gov.uk/doc/open-government-licence or write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gsi.gov.uk.

Any enquiries regarding this publication should be sent to us at: Marketing, Office of Fair Trading, Fleetbank House, 2-6 Salisbury Square, London EC4Y 8JX, or email: marketing@oft.gsi.gov.uk.

This publication is also available from our website at: www.oft.gov.uk.

CONTENTS

<u>Chapter/Annexe</u>	<u>Page</u>	
1	Executive summary	5
	Introduction	5
	Background	6
	Key effects	7
	Policy implications	7
	Application to online retailing	11
2	Introduction	13
	Research question	13
	Background	13
	Our approach	16
	Structure	19
3	Key effects on consumer surplus	21
	Introduction	21
	The distribution of gains to trade between the consumer and the firm	21
	Output	22
	The intensity of competition	25
	The ability to commit to prices	28
	Conclusion	30
4	Key characteristics	32
	Introduction	32
	Identifying key characteristics	32
	Key characteristics and consumer surplus	33
	4.a Best-response symmetry	34
	4. b The heterogeneity observed across consumers	47
	4c Dynamic interactions between firms and consumers	61
	4.d Personalised advertising	69
	4.e Costs associated with price discrimination	74

4.f	Consumer sophistication	78
4.g	Trust in online markets	83
5	Conclusions and application to online markets	88
	Introduction	88
	Framework	88
	Characteristics outside the framework	93
	Conclusion	94
6	Bibliography	95

1 EXECUTIVE SUMMARY

Introduction

1.1 On 17 May 2013 the OFT published a report on [personalised pricing online](#). As part of this work, we have undertaken a review of the economic literature regarding the following research question:

- Under what circumstances is online personalised pricing likely to cause economic harm to consumers?

1.2 Our headline conclusions are as follows:

- The general effect of personalised pricing on consumer welfare can be positive or negative, meaning that a case by case analysis is appropriate.
- Although the literature is diverse, it is possible to establish a framework to guide a case by case analysis, which indicates that:
 - Personalised pricing by a monopolist is likely to be harmful to consumer welfare, particularly where the form of discrimination is relatively sophisticated.
 - Personalised pricing where there are at least two competing firms is often not harmful to consumers, and can be beneficial.
 - While personalised pricing with competition is often not harmful, this is not always the case: in particular, it is more likely to be harmful where (i) the form of discrimination is particularly complex or opaque to consumers, (ii) it is very costly to firms and/ or (iii) consumers lose trust in the market as a result of concerns about discrimination.
 - The factors that make price discrimination more likely to be harmful are more likely to occur online than in traditional offline retail markets.

1.3 In the remainder of this section, we set out further detail regarding how we reached these conclusions. We briefly set out some background to our report. We then describe four key effects that we have identified in the literature. Finally, we describe the policy implications of our work, including the framework we have developed for the assessment of price discrimination.

Background

- 1.4 In economic terms,¹ a firm price discriminates when it charges different customers different prices that are not justified by differences in cost. Personalised pricing is a relatively refined form of price discrimination where the firm observes some heterogeneity² among consumers, and bases the price it charges on that heterogeneity.
- 1.5 For example, an online retailer may be able to observe the incomes of individual consumers on its website (or likely incomes, based on, for example, the relative affluence of the area in which the consumers live). The retailer may then choose to charge high income consumers a higher price than low income consumers.
- 1.6 We are interested specifically in price discrimination online because:
- Online retailing is an important part of the retail market (worth £68 billion a year).³
 - Online retailing is likely to allow firms to observe certain characteristics about consumers that offline retailers cannot. For example, whilst an offline retailers may be able to draw inferences about a consumer's likely income from his appearance, an online retailer may be able to observe the customer's past purchasing history, his propensity to shop around and the keywords he uses in his searches.
 - The price setting process for online sales may be different from offline retailing: for example, (i) there is usually no opportunity for the buyer to negotiate, whereas the seller is able to go back with reduced offers and (ii) price discrimination online may not be transparent or expected by the consumer in the same way as when it occurs offline.
 - There is evidence that consumers strongly object to price discrimination online.⁴

¹ Price discrimination also has a specific definition under EU law. We do not consider this, or any other legal issue, in this report.

² Heterogeneity as an economic term means diversity in or variation between consumers.

³ OFT (2013).

⁴ OFT (2013).

Key effects

1.7 The literature on price discrimination is large and diverse, with a broad range of conclusions and policy implications. Following our review of this literature, we have identified four key effects that, in combination, underpin most of the key models of price discrimination:

- Price discrimination can allow firms to charge consumers with a high willingness to pay a higher price, meaning that the firm is better off and the high willingness to pay consumers are worse off (the Appropriation Effect).
- Price discrimination can allow firms to charge consumers with low willingness to pay a lower price, meaning that more of them would buy than under a uniform price, expanding the output sold (the Output Expansion Effect).
- Price discrimination can trigger increased competition between competing firms, as each firm is better able to target customers of other firms with a competitive price offer (Intensified Competition Effect).⁵
- Price discrimination can prevent firms from committing not to decrease their prices in the future (the Commitment Effect).

1.8 All of the effects, with the exception of the Appropriation Effect, are such that price discrimination actually has a positive effect on consumer surplus. In practice, the overall effect of price discrimination will depend on how these four effects balance.

Policy implications

1.9 By way of policy implications, we have created a simple framework that indicates under what circumstances price discrimination is likely to be

⁵ For example, if a particular consumer has a preference for firm Y over firm X, firm X would like to charge that consumer a low price, as it knows that otherwise that consumer will purchase from firm Y. If the firm can only charge a uniform price, then charging that particular consumer a low price would require the firm to decrease the price it offers to all consumers. If, on the other hand, the firm is able to price discriminate, then it can offer that particular consumer a low price without having to decrease its prices to all consumers. In other words, with price discrimination the firm is incentivised to offer that consumer a lower price than with uniform pricing. Price discrimination effectively allows firms to compete for each consumer individually, which intensifies competition between the two firms and generally results in lower prices. Importantly, this is true for all consumers, as exactly the same logic applies to a consumer that prefers firm X over firm Y.

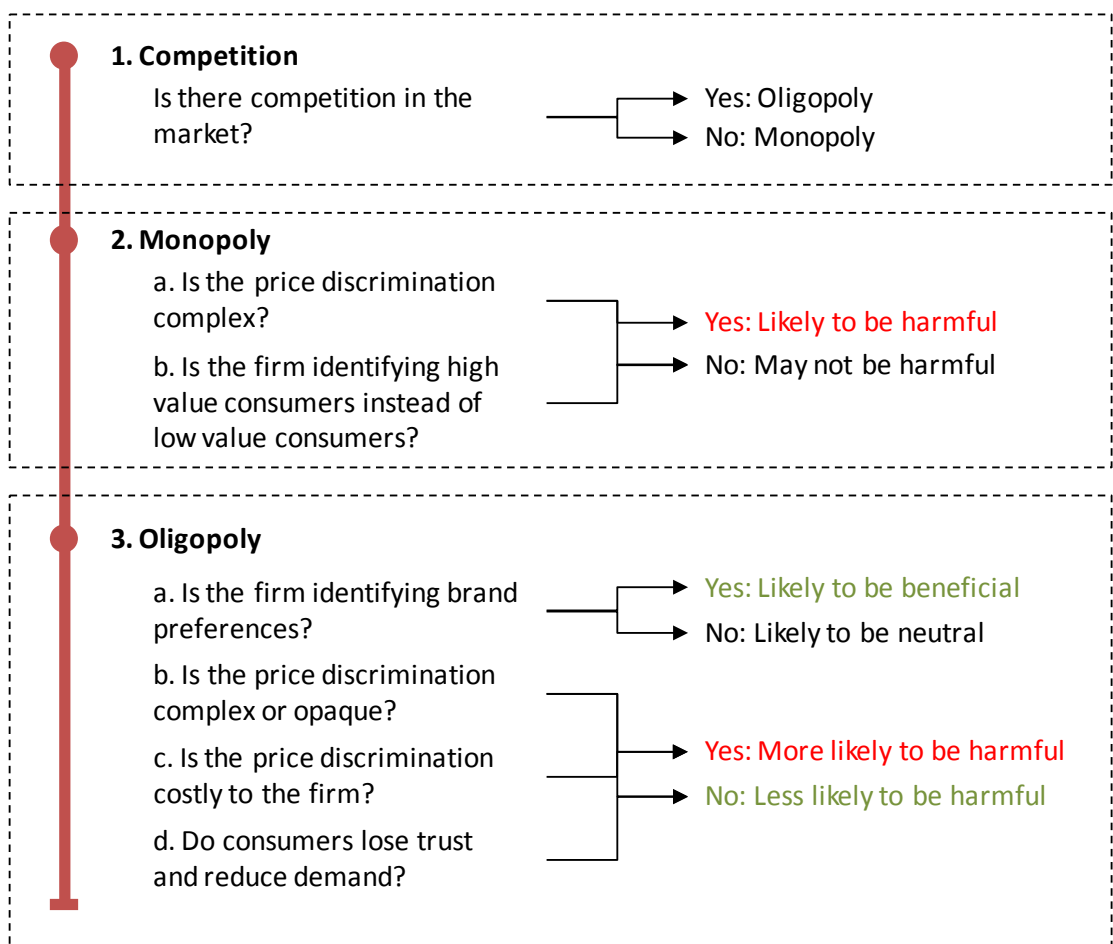
harmful in economic terms to consumers. This framework could guide the analysis in an investigation of a particular case of price discrimination.⁶

1.10 We describe the framework, before describing its application to online markets.

Framework

1.11 We summarise our framework in the figure below, the first step of which is to identify whether there is a monopoly or an oligopoly (if there is perfect competition, then price discrimination will not occur).

Figure 1: Summary of framework



Monopoly

⁶ This framework describes when price discrimination will improve aggregate consumer surplus. In many cases, there may be individual winners and losers, such that some consumers are made worse off and some are made better off. Even if aggregate consumer surplus increases, the OFT may be concerned if vulnerable consumers are made worse off.

1.12 Where there is only one firm in the market then the Intensified Competition Effect does not apply,⁷ and the overall effect of price discrimination typically depends on how the Output Expansion Effect balances against the Appropriation Effect. The literature notes that the latter will dominate the former and price discrimination will be harmful if:

- The monopolist is identifying and segmenting high value consumers instead of low value consumers.⁸
- The form of discrimination is relatively complex, such that the monopolist is identifying a significant amount of heterogeneity between consumers that allows it to segment consumers into numerous groups.⁹ Given that we define personalised pricing as sophisticated direct price discrimination, in certain cases it may be sufficiently complex to reduce consumer surplus.

Oligopoly

1.13 In an oligopoly model, the Intensified Competition Effect can cause price discrimination to have a positive effect on consumers.

1.14 Consider, for example, if an online retailer was able to identify whether a particular consumer had purchased from a competing online retailer. Price discrimination would allow that retailer to charge that consumer a lower price in an attempt to steal his business from his competitor. This form of increased competition occurs when firms are able to observe or infer some kind of brand preference on the part of their consumers.

1.15 If oligopolists observe something unrelated to brand preference (such as switching costs, or the extent to which the consumer shops around), then the Intensified Competition Effect does not operate. Instead, price discrimination will create winners and losers among consumers, but the aggregate effect on consumer surplus will be neutral.

1.16 On the face of it, this creates a relatively strong conclusion regarding price discrimination with competition: it is at worst neutral, and at best positive.

⁷ Note that we are referring to a monopoly in the economic sense, which is distinct from the legal question of whether a firm is dominant under EU law.

⁸ If the monopolist is merely segmenting its highest value consumers then there may not be much of an Output Expansion Effect. If the monopolist is segmenting low value consumers, then it is specifically charging those consumers a low price and therefore likely to price them into the market.

⁹ As the monopolist segments consumers into finer and finer groups, the potential for output to increase becomes limited, as total output is clearly limited by the number of prospective consumers for that product.

1.17 This is not always the case, however: we have identified the following four general features that, if present in an oligopoly market, may mean that price discrimination is harmful (in each case we state the result and present some intuition in a footnote):

- Consumers are not sophisticated, in that they do not account for how their purchasing decisions today affect the prices they are offered tomorrow.¹⁰
- Price discrimination is not transparent to the consumer.¹¹
- It is very costly to firms to engage in price discrimination.¹²
- Consumers lose significant trust in online markets due to concerns about price discrimination.¹³

1.18 The presence of any of these features does not mean that price discrimination is necessarily harmful to consumers. For example, price discrimination that is not transparent to the consumer is not necessarily harmful to the consumer, as in some cases competition between firms will benefit consumers even if they do not know about it. This framework is intended only as a guide for investigations into specific cases of price discrimination.

1.19 We summarise this framework in the following table, in which we discuss each of the major issues, as well as list some observable features that may be important in an investigation.

Table 1: Summary of framework

Issue	Observable feature	Effect of price discrimination on consumer surplus
1. Level of competition	- number of firms in market - barriers to entry.	Go to 2. or 3. below

¹⁰ In some cases, consumers will be better off if they are able to understand that price discrimination is occurring and respond accordingly. For example, a sophisticated consumer may search extensively online for cheap deals, and thus build up an online profile as a savvy shopper and possibly get charged lower prices. This is not possible, however, if consumers do not understand how price discrimination is occurring, in which case consumers may be worse off.

¹¹ The intuition to this is analogous to the previous footnote: if consumers are not aware that price discrimination is occurring, then clearly they may not respond optimally in certain cases.

¹² If firms incur significant costs to price discriminate (the cost of information collection and analysis, for example), then they will need to recover these costs through higher prices.

¹³ If consumers lose trust in online markets, then they may reduce their demand, which causes potential gains to trade to go unexploited.

2. Monopoly	- complexity of discrimination (number of groups segmented) - type of consumer being identified.	Likely to be harmful if sophisticated or identifying high value consumers
3. Oligopoly	- identifying brand preferences - identifying some other form of heterogeneity.	Positive if identifying brand preferences, neutral otherwise, subject to other issues below
3.a. Sophistication of consumer	- complexity of price discrimination - transparency.	More likely to be harmful if complex or not transparent
3.b. Cost of price discrimination	- source of information being collected - cost to the firm of collection.	More likely to be harmful if costly
3.c. Loss of market trust	- reduction in demand caused by introduction of price discrimination - consumer response to price discrimination - transparency of price discrimination.	More likely to be harmful if reduces trust in market and demand.

Application to online retailing

1.20 Online personalised pricing is likely to take various different forms depending on, among other things, (i) the level of competition in the market, (ii) the complexity of price discrimination (and thus how likely consumers are to understand it) and (iii) the cost to the firm of discriminating. It is not possible, therefore, to conclude whether, in general, online personalised pricing is harmful or beneficial to consumers.

1.21 Instead, in specific investigations it will be necessary to follow the framework we summarise above and discuss in detail in the remainder of this report, based on the specific factors of the case.

1.22 We note, however, that many of the market characteristics that make price discrimination more likely to be harmful apply to online retailing more than they do to offline retailing. In particular, online retailing is (i) more likely to be complex, as firms can observe more heterogeneity and (ii) more likely

to be opaque, as consumers may not be able to observe the price being quoted to other consumers. This does not mean that personalised pricing online, as a rule, is likely be harmful: it just means that caution should be taken when applying this framework to individual cases of price discrimination.

2 INTRODUCTION

Research question

2.1 The OFT launched a call for information on online personalised pricing practices on 15 November 2012. As part of this, we have undertaken a review of the relevant economic literature, with a focus on the following research question:

- Under what circumstances is personalised pricing online likely to cause economic harm to consumers?

Background

2.2 By way of background, we first introduce the economic concept of price discrimination and personalised pricing. We then consider how these apply to online markets.

Price discrimination

2.3 In general terms, price discrimination occurs where there are differences in the price of similar products that are not justified by differences in cost.

Common examples include:

- Student discounts: students commonly receive discounts for train or theatre tickets. Since the cost of train travel or theatre attendance does not vary depending on whether the customer is a student, this is a form of price discrimination.
- Airline tickets: airline tickets are commonly more expensive if purchased close to the date of departure, even though the cost of an additional passenger does not depend on whether that passenger purchased his ticket early or late.

2.4 More formally, Armstrong (2006) defines price discrimination as the situation where two 'similar' products with the same marginal cost are sold by a firm at different prices.¹⁴

¹⁴ Armstrong (2006), page 97.

- 2.5 Alternatively, Stigler (1987) proposes a more general definition of price discrimination as the situation in which two 'similar' products are sold at prices that are in different ratios to their marginal costs.¹⁵
- 2.6 Price discrimination can be direct or indirect. Direct price discrimination occurs where the firm observes some form of heterogeneity among its customers and is able to vary its price based on that heterogeneity. Indirect price discrimination occurs where the firm cannot observe the heterogeneity in its customers, but is able to offer a menu of options that results in different consumers self-selecting different options according to their preferences.
- 2.7 For example:
- Student discounts are direct, in that the firm can observe whether a consumer is a student and then condition its prices on that fact.
 - Quantity discounts (such as 'buy three and get 10 per cent off') are indirect, as consumers will self-select different price and quantity combinations, thus paying a different overall price.
- 2.8 The economic literature frequently makes reference to first-, second- and third-degree price discrimination (Pigou (1932)):
- First-degree price discrimination occurs where the firm is able to perfectly observe all the relevant¹⁶ heterogeneity among consumers and price discriminate accordingly (it can therefore be thought of as perfect direct price discrimination).
 - Second-degree price discrimination occurs as a form of indirect price discrimination, where the firm offers a menu of options.
 - Third-degree price discrimination is imperfect direct price discrimination, in that the firm observes some but not all of the relevant heterogeneity among consumers.

¹⁵ Stigler (1987). Strictly, therefore, charging the same price on two similar goods when they have different marginal costs is also price discrimination.

¹⁶ *Relevant* heterogeneity in this case means heterogeneity that affects the willingness to pay of consumers for that particular good.

Personalised pricing

- 2.9 The OFT's call for information concerns 'personalised pricing' online. In economic terms, we take personalised pricing to mean direct price discrimination that is imperfect but sophisticated. In other words, for the purposes of this report we take personalised pricing to mean sophisticated third-degree price discrimination.
- 2.10 Perfect direct price discrimination and indirect price discrimination are not included within the scope of this review, although we discuss them below where informative.
- 2.11 In the remainder of this report we use the economic term price discrimination. Unless explicitly stated otherwise, we are referring to imperfect, direct price discrimination.

Application to online markets

- 2.12 Price discrimination occurs in many 'traditional' offline markets. Any transaction that involves some form of negotiation between the buyer and the seller (which includes almost all business to business transactions, as well as home improvement) is likely to involve an informal assessment by the seller of the buyer's willingness to pay ('WTP').
- 2.13 We are interested specifically in price discrimination online because:
- Online retailing is an important part of the retail market (worth £68 billion a year).¹⁷
 - Online retailing is likely to allow firms to observe certain characteristics about consumers that offline retailers cannot. For example, whilst an offline retailers may be able to draw inferences about a consumer's likely income from his appearance, an online retailer may be able to observe the customer's past purchasing history, his propensity to shop around and the keywords he uses in his emails.
 - The price setting process for online sales may be different from offline retailing: for example, (i) there is usually no opportunity for the buyer to negotiate, whereas the seller is able to go back with reduced offers and

¹⁷ OFT (2013).

(ii) price discrimination online may not be transparent or expected by the consumer in the same way as when it occurs offline.

- There is evidence that consumers feel strongly against price discrimination online.¹⁸

2.14 In short, therefore, we are interested in online pricing because the market is very large, and there is some evidence that price discrimination online has particular features that are worthy of further investigation. We discuss the application to online markets in detail in a later section.

Our approach

2.15 We first discuss three particular aspects of our approach regarding (i) measuring consumer harm as an economic concept, (ii) the conditions for personalised pricing to be feasible, (iii) exploitative and exclusionary conduct and (iv) consumer law and privacy regulations.

2.16 We then set out a framework for analysis more generally, and discuss how we synthesise the literature to reach general conclusions.

Measuring harm

2.17 Our research question concerns ‘economic harm’ to consumers. We measure the benefit of purchasing a product in economic terms as consumer surplus, which is the difference between what the consumer paid and what he would have been willing to pay. Harm, as an economic concept, is a decrease in consumer surplus.

2.18 Harm to consumers could be interpreted as including other, non-economic factors, such as anxiety and distress. We briefly consider these non-economic factors to the extent they affect consumer surplus, but we do not consider them otherwise.

2.19 Where relevant we also discuss the effect on profits and thus total surplus (which is the sum of consumer surplus and profits).

¹⁸ OFT (2013).

The feasibility of personalised pricing

2.20 Certain conditions must be met for a firm to be able to feasibly engage in personalised pricing:¹⁹

- The firm must have some degree of market power, as if the market is perfectly competitive then the law of one price holds, as competition drives the price for every consumer down to marginal cost.
- The firm must be able identify some heterogeneity among consumers, and condition prices on that heterogeneity.
- Arbitrage must not be possible: that is, consumers must not be able to resell the good easily and at low additional cost, as otherwise arbitrage would result in a uniform effective price.

2.21 We have not sought to assess the feasibility of price discrimination online in general, although we note that responses to the call for information did not indicate that price discrimination was generally infeasible. In any investigation of price discrimination its existence (and implied feasibility) should be observable.

2.22 In the remainder of this report, we assume that price discrimination is feasible unless we explicitly state otherwise.

Exploitative and exclusionary conduct

2.23 Price discrimination can be considered as potentially an exploitative conduct, in that firms may use market power to exploit consumers and appropriate their consumer surplus.

2.24 However, in some cases it could also be used as an exclusionary device: if a dominant undertaking offers selective prices to potential customers of smaller rivals for the purpose of excluding or deterring competition this may constitute an abuse of a dominant position under the Chapter II Prohibition /Article 102 of TFEU (for example. Akzo Chemie BV v Commission and the Irish Sugar case.)

¹⁹ See Stole (2007), page 3.

2.25 The focus of the call for information, however, is on exploitative concerns rather than exclusionary issues. As a result, we do not consider exclusionary issues further below.

Consumer law and privacy regulations

2.26 Price discrimination based on observed heterogeneity is also subject to consumer law and privacy regulations, including but not limited to the Privacy and Electronic Communications (EC Directive) Regulations 2003, the Data Protection Act 1998 and the Consumer Protection from Unfair Trading Regulations 2008. Additional detail is contained in OFT (2013).

2.27 This report is limited to the economic analysis of the effect of price discrimination on consumer surplus, and does not consider consumer law or privacy regulations.

A framework for our analysis

2.28 There are numerous models of price discrimination in the literature, each with differing assumptions and implications for consumer surplus. It is not possible to select a single model as the best representation of online retailing, as online retailing can vary significantly.

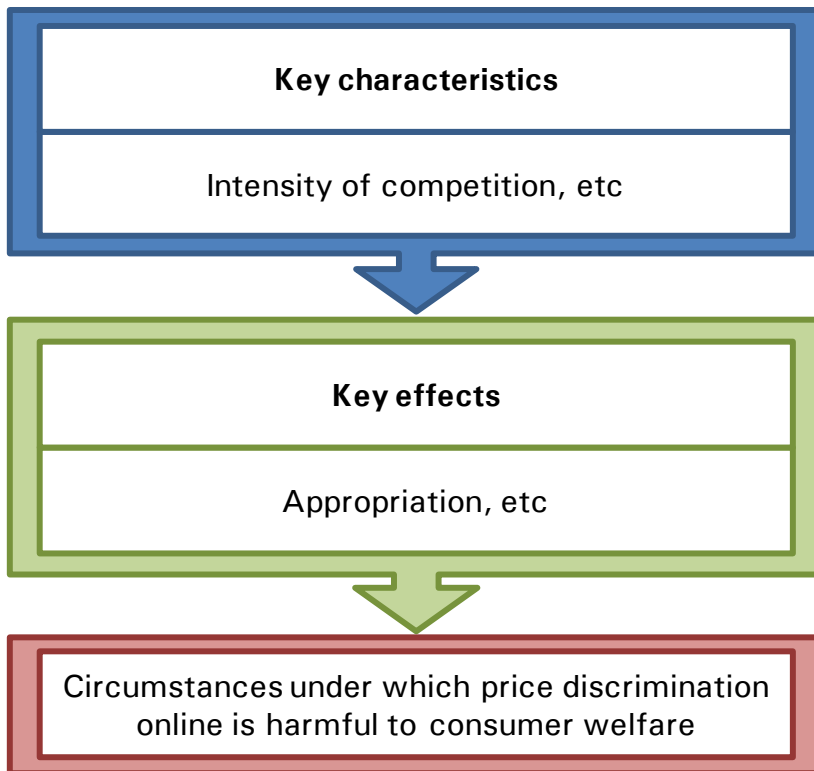
2.29 It is, however, possible to identify various effects that are common across models. For example, one effect we have identified is that price discrimination allows the firm to appropriate some, or all, of the consumer surplus.

2.30 These effects are driven by the underlying assumptions of each model (that is, the key characteristics that differ between the models). For example, a key characteristic is the intensity of competition in the market.

2.31 The overall effect of price discrimination on consumer surplus depends on the relative balance of the key effects. If, for example, the key characteristics are such that the dominant effect is the appropriation of consumer surplus, then clearly consumer surplus is reduced by price discrimination. On the other hand, if the key characteristics are such that price discrimination results in output increasing significantly, then it may be that consumer surplus is increased.

2.32 Figure 2 below summarises our general approach, based on identifying key effects and key characteristics, and using this to create a framework to understand when price discrimination may or may not be harmful.

Figure 2: A framework for reviewing the literature



2.33 The aim of this research is not to establish definitively whether price discrimination is harmful to consumer surplus in general, as this will vary depending on the context. Instead, we use the framework set out above to identify the circumstances under which price discrimination may be harmful to consumer surplus.

Structure

2.34 The remainder of this report is structured as follows:

- Section 3: we introduce the key effects of price discrimination on consumer surplus.
- Section 4: we introduce the key characteristics in the literature that drive these effects, before considering each of the key characteristics individually in Section 4.a to 4.g.

- Section 5: we set out our conclusions and apply the findings from our review specifically to online markets.

3 KEY EFFECTS ON CONSUMER SURPLUS

Introduction

3.1 In this section, we describe the key effects that we have identified in the literature. We have identified four key effects, which we describe by starting with a simple static model before building up to a more complex model:

- We first set out a static model of monopoly pricing, and specifically examine the effect of price discrimination:
 - on the distribution of gains to trade between the consumer and the firm
 - giving rise to any expansion in output.
- We then introduce competition, and consider the effect of price discrimination on the intensity of competition.
- We finally introduce a dynamic aspect to interaction between firms and consumers, in order to consider the ability of the firm to commit to prices.

3.2 We describe each of the key effects below in turn. Our focus, as discussed above, is on consumer surplus, but we discuss firm profits and total surplus where relevant.


3.3 We use these key effects to introduce some of the common economic issues that underpin the broad range of literature on price discrimination. Note that whilst these are, in our opinion, the most important effects, they are not the only effects that feature in the literature, as we will discuss in Section 4.

The distribution of gains to trade between the consumer and the firm

3.4 First, consider a simple model of price discrimination under monopoly. The gains to trade are the differences between each consumer's WTP and the marginal cost ('MC') to the firm of producing the good. Under uniform pricing, these gains to trade are shared between the consumers and the

firm, who receive consumer surplus and producer surplus (or profit), respectively, as shown in the figure below.²⁰

Figure 3: Consumer surplus under a monopoly with uniform pricing



A: Consumer surplus
B: Producer surplus
C: Deadweight loss

- 3.5 If the monopolist is then able to price discriminate, it will set different prices for different groups of consumers (at the extreme, under first-degree price discrimination it is able to charge every consumer a different price). This results in the 'Appropriation Effect', whereby the monopolist identifies groups of consumers with high willingness to pay and charges them a higher price, thereby appropriating some of their consumer surplus.
- 3.6 That is, some or all of the existing consumer surplus (area A in Figure 2) is converted into producer surplus, meaning the Appropriation Effect reduces consumer surplus and increases producer surplus.


Output

- 3.7 Even under a simple static monopoly model the effect of price discrimination on consumer surplus is not limited to the Appropriation Effect. Consider a consumer whose WTP is just below the uniform price P_0 (that is, consumer Q_0+1). The monopolist has an incentive to sell to that consumer, as his WTP exceeds marginal cost, so there are potential gains to trade.

²⁰ Deadweight loss, as presented in the figure, relates to potential gains to trade that are unexploited.

- 3.8 Under uniform pricing, however, the firm sets a single price for all consumers. Reducing the price so as to sell to consumer Q_{0+1} would mean reducing the price to all consumers (who, by definition, have a higher WTP) and reducing overall profits. In general, a monopolist with uniform pricing will set that price above MC, thus leaving some gains to trade unexploited. This is the deadweight loss triangle set out in Figure 2 above.
- 3.9 Consider, however, if the monopolist can price discriminate. In this case, the monopolist could profitably reduce the price it charges to consumer Q_{0+1} without affecting the price it charges to other consumers. That is, price discrimination allows the monopolist to profitably sell greater volumes ('the Output Expansion Effect'). Depending on the type of price discrimination, there may then be some net gain to consumer surplus resulting from these additional sales.
- 3.10 For example, rail firms commonly price discriminate between students and non-students. Charging non-students an increased price allows firms to charge the students a lower price, thus allowing more students to travel than would do without price discrimination. This may increase aggregate consumer surplus.
- 3.11 The figure below illustrates the countervailing Appropriation Effect and Output Expansion Effect in a static model of monopoly pricing:
- Figure 4a shows the division of gains to trade under uniform pricing.
 - Figure 4b demonstrates the division of gains to trade assuming the monopolist can third-degree price discriminate between consumers with low and high WTP.
 - Figure 4c demonstrates the changes resulting from price discrimination.

Figure 4: Static model of monopoly price discrimination



3.12 The ability to price discriminate allows the monopolist to appropriate some consumer surplus (D). However, price discrimination also allows those consumers between Q_0 to Q_{1L} that were previously not purchasing to purchase, with associated consumer and producer surplus (G and H).

3.13 This example is based on third-degree price discrimination where the monopolist can segment the market into two groups. If the monopolist can perfectly price discriminate then (i) the Appropriation Effect will convert all consumer surplus into producer surplus and (ii) the Output Expansion Effect has no impact on consumer surplus, as all the gains to trade on the additional sales are appropriated by the firm.

3.14 This simple illustration raises some important themes that we will explore in more detail in the next section:

- The effect of price discrimination on consumer surplus is ambiguous (for example, in the model above, D and E of consumer surplus is lost, but H is gained: based on this graphical presentation it appears that consumer surplus decreases, but this is not generally the case).
- Price discrimination may increase total surplus (that is, the deadweight loss may be reduced), even if consumer surplus might decrease.

- Price discrimination usually results in winners and losers (for example, in the model above consumers 0 to Q0 are worse off and consumers Q0 to Q1L are better off).

The intensity of competition

3.15 If there is more than one firm in the market, the ability to price discriminate could affect the intensity of competition between the firms. We first set out the intuition for this result, before presenting a simple economic model.

Intuition

- 3.16 Consider, for example, two firms each selling a differentiated product. Some consumers prefer firm A's product, and some prefer firm B's product. At equal prices, consumers will always select their preferred product. However, if firm A's prices are higher, some A-type consumers may instead select firm B's product (their least-preferred option).
- 3.17 Under uniform pricing, each firm sets a single price for all consumers, regardless of their brand preferences. If the firm is able to observe each customer's brand preference and can price discriminate, then it can charge different consumers different prices depending on their brand preference. This results in the two firms competing more intensely, because the firms can cut prices for certain consumers without that affecting the prices they charge other consumers.
- 3.18 Consider, for example, a consumer that prefers brand A. Under uniform pricing, firm B would have to offer that consumer the same price as every other consumer. If price discrimination is feasible, then firm B can observe that the consumer prefers brand A (and thus has a lower WTP for brand B) and therefore it can charge a lower price without lowering the price it charges to consumers that prefer B. We refer to this as the 'Intensified Competition Effect' (this is also known in the literature as the 'business-stealing effect'²¹ or the 'customer poaching effect').²²
- 3.19 The intuition for this result is the same as above for the Output Expansion Effect: price discrimination allows firms to charge low WTP consumers a lower price, thus allowing them to compete more intensely for each other's customers, without affecting the price they charge other customers. In

²¹ Stole (2006), page 19.


²² Armstrong and Vickers (2011), page 6.

other words, price discrimination means each consumer is in a separate market to be contested individually.

A simple economic model of the Intensified Competition Effect

3.20 This effect can also be shown more formally, based on Thisse and Vives' (1988) Hotelling model.

Figure 5: Hotelling model



3.21 Two firms are situated at each end of the unit interval [0, 1], along which consumers are uniformly located. A consumer located at x (somewhere on the unit interval) is a distance $d_A=x$ from firm A and $d_B=1-x$ from firm B. Distance in a Hotelling model such as this can be interpreted as representing consumers' preferences over two differentiated products or over two differentiated online retailers.

3.22 Suppose a consumer wishes to buy a single unit from either firm A or firm B. If he buys from firm $i = A, B$ his net surplus is:

$$u_i = v - p_i - 0.5d_i$$

3.23 v signifies the underlying value the consumer places on the unit, irrespective of the firm from it is purchased,²³ P signifies the price charged by firm i and marginal cost is normalised to zero. d measures the differentiation between A and B. The value 0.5 simply captures how much utility the consumer loses from having to choose his least preferred option.²⁴

3.24 A consumer is in firm A's strong segment if $x < 0.5$ and in firm B's strong segment if $x > 0.5$. For those consumers in its strong segment, firm A will set its price such that firm B cannot profitably serve that consumer. That is, P_A is set such that the consumer will prefer firm A even if $P_B=0$.

²³ That is, differentiation between the firms is captured by their location, instead of differences in the underlying value, v .

²⁴ We choose 0.5 for this value for ease of exposition, as it does not affect the general conclusions. We discuss a more general model where this parameter is allowed to vary in Section 4.b below.

$$\begin{aligned}
 &U_A > U_B \\
 &v - P_A - 0.5x > v - P_B - 0.5(1 - x) \\
 &P_A < 0.5 - x
 \end{aligned}$$

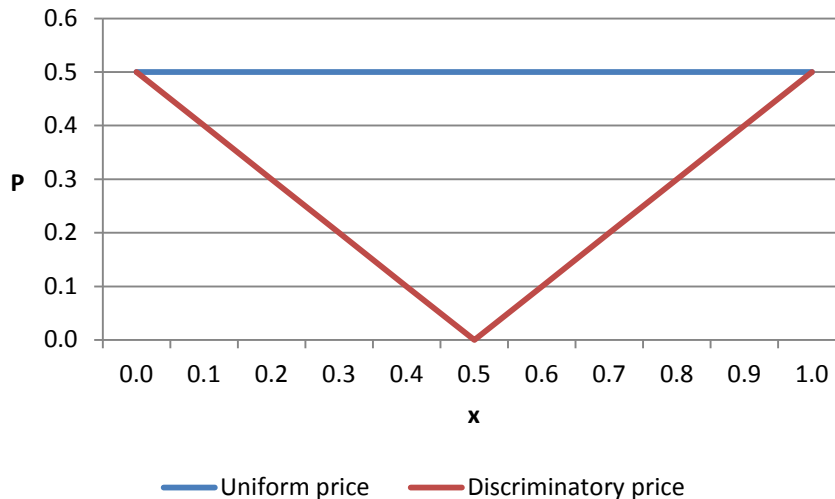
3.25 Solving for equilibrium price:

$$P_x = \{0.5 - x \text{ if } x \leq 0.5 | x - 0.5 \text{ if } x \geq 0.5\}$$

3.26 It is straightforward to show that the uniform price is $P=0.5$ (intuitively, each firm makes the marginal consumer (that is, the consumer at $x=0.5$) indifferent between the two goods). The uniform price is clearly greater than the discriminatory equilibrium price for every x . The benefit of price discrimination is greatest for the consumer in the middle.

3.27 The figure below shows the uniform and discriminatory prices that would be charged to each consumer.

Figure 6: The Increasing Competition Effect – Prices



3.28 In the figure below we set out unit consumer surplus.²⁵

Figure 7: The Increasing Competition Effect – Consumer surplus

²⁵ For the purposes of this illustration we assume that consumers have the same intrinsic valuation for the good, $v=1$, and differ only in their brand preferences.



3.29 Although all consumers are better off under price discrimination (aside from those exactly at either end, who are as well off), some consumers benefit more than others. Consumers in the middle of the range (that is, consumers that are largely indifferent between the two goods) are worst off under uniform pricing but best off under price discrimination.

3.30 We have used this model to introduce the concept of an Intensified Competition Effect. It is important to note, however, that this model makes various assumptions about, for example, the form of demand and the distribution of consumers. We consider a more general model in Section 4.a below.

The ability to commit to prices

3.31 The models and intuition discussed thus far have been static. If the firm(s) and the consumer interact more than once; that is, the interaction becomes dynamic. We extend the intuition to a dynamic game across two periods between consumers and a single firm. Consumers may buy in all, one or none of the periods.

3.32 In a dynamic game the consumer can choose when to purchase the product. The monopolist can observe whether the consumer previously bought from it. If the consumer did not buy from it in the previous period, then that is a signal that the consumer has a low WTP (the consumer did not purchase because his WTP was lower than the price being offered). If

the consumer did purchase in the previous period, then conversely that is a signal that the consumer has a high WTP.

- 3.33 If the monopolist can price discriminate, it will charge a uniform price in the first period. In the second period, it will charge a low price to consumers with a low WTP (non-returning customers) and a high price to consumers with a high WTP (returning customers).
- 3.34 If consumers are sophisticated, however, then they take the monopolist's likely actions in the future into account in the present.²⁶ That is, they know that by not purchasing in the first period, even if they are high WTP, then they will receive a lower price in the second period, because they are incorrectly identified as low WTP. In other words, the sophisticated high WTP consumer can pretend to be a low WTP consumer.
- 3.35 We do not go through the formal model here (see Armstrong (2006)) but, in simple terms, the high WTP consumer is able to earn additional consumer surplus, because he can pretend to be low WTP. The intuition is that consumers have more information than firms (because consumers know their WTP, whereas firms know only if they are returning or non-returning customers), and thus extract an informational rent from the firm. Firms earn a lower profit, and would be better off if they did not price discriminate.
- 3.36 As a result, the firm may wish not to price discriminate. However, even if it claims in the first period that it will not price discriminate in the second period (meaning consumers have no incentive to pretend to be low WTP), once the monopolist reaches the second period it will find that it is, in fact, optimal to price discriminate to take advantage of the high WTP consumers that revealed themselves in the first period. The sophisticated consumer knows that the monopolist's claim is not credible, and thus does not change his behaviour.
- 3.37 In other words, the firm faces a commitment problem. This is analogous to the commitment problem that monopolists face when selling durable consumer goods, as first shown by Coase (1972): the consumer knows the monopolist will always have an incentive to drop the price tomorrow so will not purchase today.

²⁶ We discuss whether this is a realistic assumption in Section 4.

3.38 Note that this commitment problem results from the ability to price discriminate. If the monopolist was unable to price discriminate between returning and non-returning customers, there would be no incentive for even sophisticated consumers to pretend to be low WTP. Banning price discrimination would, in this context, restore the monopolist's commitment power to the detriment of consumers.

3.39 This is one specific model: there are many models of price discrimination in dynamic contexts, some of which have different conclusions, as we discuss in Section 4 below. The Commitment Effect, however, is a recurring theme within these dynamic models.

Conclusion

3.40 We have introduced the four main effects we have identified. Many of the models in the literature can be expressed in terms of combinations of these effects.

3.41 We have also shown that three out of the four effects affect consumer surplus positively, as summarised in Table 1 below. Two of the four affect producer surplus positively.

Table 2: Direction of the effects on consumer surplus

Effect	Moving from uniform pricing to discriminatory pricing:	
	Consumer surplus	Producer surplus
Appropriation Effect	Negative	Positive
Output Expansion Effect	Positive	Positive
Intensified Competition Effect	Positive	Negative
Commitment Effect	Positive	Negative

3.42 In the course of illustrating these effects we have also touched upon various model characteristics that determine the effect of price discrimination on consumer surplus, including:

- the level of competition in the market

- whether the interaction between consumers and firms is static or dynamic, and
- the level of sophistication of consumers.

3.43 We discuss these and other key characteristics in greater depth in the next section.

4 KEY CHARACTERISTICS

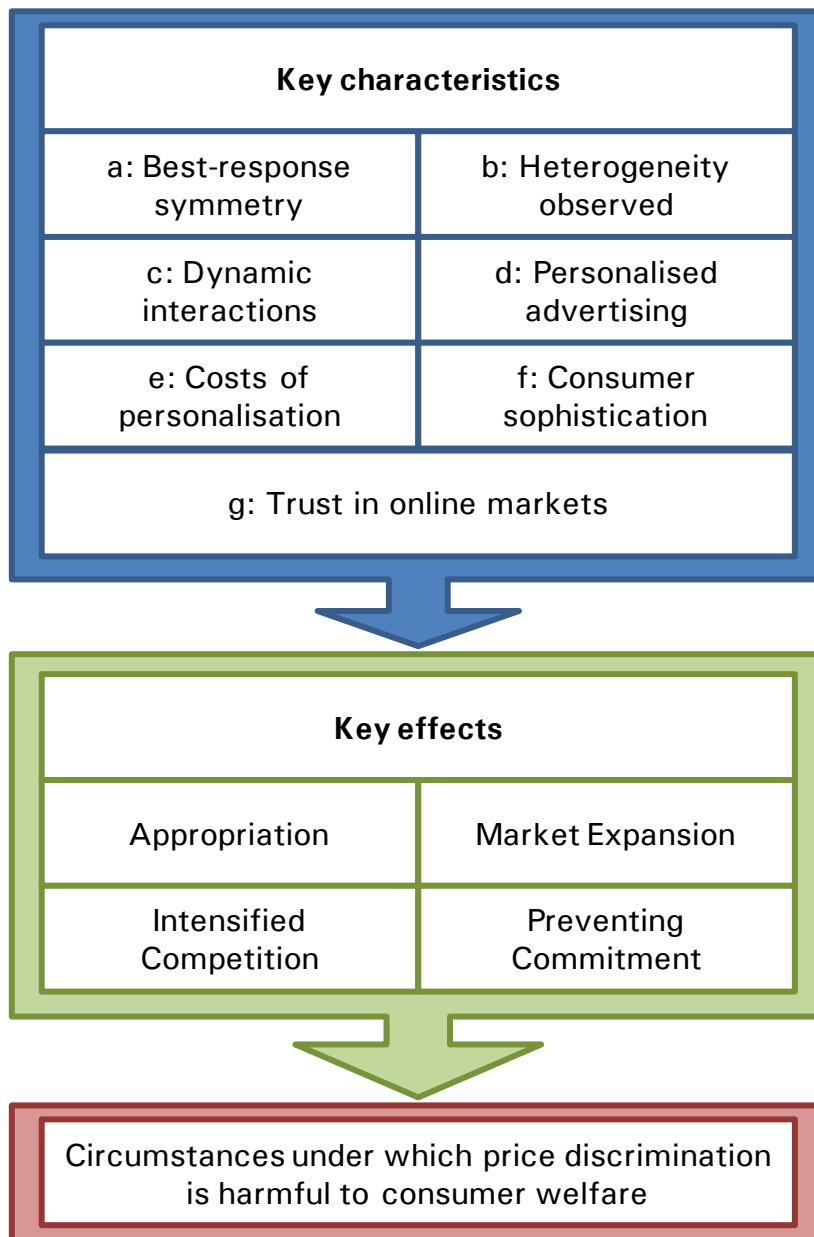
Introduction

- 4.1 In this section, we introduce the key characteristics we have identified from the literature. By key characteristics we mean the ways in which the application of price discrimination can vary.
- 4.2 We then discuss two issues that affect the link in the literature between these characteristics and consumer surplus: (i) interactions between key characteristics and (ii) the effect of assuming unit demand.

Identifying key characteristics

- 4.3 Based on our review of the literature, we have identified the following relevant characteristics, which we discuss in turn in the following sections:
- Section 4.a: best-response symmetry (the extent to which firms agree or disagree on their ranking of consumers)
 - Section 4.b: the type, amount and accuracy of the heterogeneity observed by the firm(s)
 - Section 4.c: dynamic interactions between firms and consumers
 - Section 4.d: the interaction of personalised pricing with personalised advertising
 - Section 4.e: costs associated with personalised pricing
 - Section 4.f: the level of consumer sophistication
 - Section 4.g: the level of trust in online markets.
- 4.4 Whilst the models in the literature may vary in additional ways, these are the variations we consider to be particularly important. Note that this list also covers characteristics that are not considered in detail in the literature, but that we consider are important for considering how theoretical models of price discrimination can be applied to online pricing in practice.

Figure 8: Our approach



Key characteristics and consumer surplus

- 4.5 In the following sections, we review the main conclusions from the literature on each of the characteristics. There are two issues that should be borne in mind when considering the link in the literature between these characteristics and consumer surplus.

- 4.6 First, there are interactions between these characteristics that may be difficult to disentangle. In particular, the level of competition has implications for many different characteristics, including for example how the type of heterogeneity affects consumer surplus. Nevertheless, we consider that a characteristic by characteristic review of the literature is most informative, provided that these interactions are made clear. We do not consider the level of competition individually, but we take it into account in our discussion of other characteristics and in our conclusions.
- 4.7 Second, the ability of the literature to identify the effect of these characteristics on total surplus is, in some cases, limited. This is particularly the case where unit demand models are used, where unit demand means that each consumer demands exactly one unit, subject to his reservation price being met.
- 4.8 The assumption of unit demand is used for its tractability, but it means that total output in the market is fixed regardless of price, when in most cases one would intuitively expect total quantity to change with price. This assumption makes it difficult to draw conclusions regarding the effect of price discrimination on total surplus.²⁷
- 4.9 However, it is still possible to observe the division of surplus between producers and consumers. It is possible, therefore, to accurately gauge whether the introduction of price discrimination would make consumers better or worse off.

4.a Best-response symmetry

Introduction

- 4.10 In the example discussed at paragraph 0 above, the monopolist is able to segment its consumers into high WTP and low WTP, or, as it is sometimes called in the literature, a 'strong' group and a 'weak' group.²⁸ The Appropriation Effect arises as a result of the monopolist charging the strong group a higher price relative to the uniform price and the weak group a lower price.

²⁷ For further discussion of this issue see Armstrong (2006), footnote 30, and Stole (2006), page 39.

²⁸ Robinson (1933).

4.11 Now consider if there are two firms in the market, along with heterogeneous consumers. An important feature of any oligopoly models of this sort is whether the two firms identify the same group of consumers as being strong, which Corts (1998) terms 'best-response symmetry', or whether one firm's strong group is the other firm's weak group ('best-response asymmetry').

4.12 For example:

- if consumers have different brand or store preferences and the firms are able to observe those differences (as in the Thisse and Vives (1988) model set out at paragraph 4.18 below), then that would likely be best-response asymmetric: a consumer that prefers firm A over firm B is in firm A's strong group and in firm B's weak group, and
- if firms are able to observe how likely consumers are to shop around, then that would likely be best-response symmetric: both firms A and B would identify consumers that do not shop around as being in their strong group (that is, they could be charged a higher price than those that do shop around).

4.13 The intuition regarding why best-response symmetry matters is as follows. Best-response asymmetry means that firms are able to identify customers that are pre-disposed to prefer the other firm, and thus charge them a lower price in an attempt to poach them. This can result in more intense competition across consumers and lower prices.

4.14 If, however, the heterogeneity that is observed among consumers does not enable firms to observe which consumers are pre-disposed to prefer the other firm, they cannot try to poach them. The effect of the observed heterogeneity in this case is more ambiguous.

4.15 We discuss best-response asymmetry and symmetry in theory, before considering it in practice.

Best-response asymmetry in theory

4.16 We first briefly describe a Hotelling model of best-response asymmetry, before discussing a more general treatment.

Hotelling model

- 4.17 We describe best-response asymmetry and present Thisse and Vives' (1988) Hotelling model in Section 2 above. Assuming a uniform distribution of consumers, intensified competition causes prices to fall for all consumers under price discrimination.²⁹
- 4.18 Price discrimination in this case also causes firm's profits to decrease: in the Thisse and Vives (1988) model profits under perfect price discrimination are half profits under uniform pricing.³⁰ Effectively, the Appropriation Effect is dominated by the Intensified Competition Effect.³¹
- 4.19 This comes about because it is optimal for each firm to price discriminate given each other's prices. Consider, for example, the situation in which firm A was able to price discriminate but firm B was not. Firm A would be able to target firm B's strong consumers, whereas firm B would not be able to target firm A's consumers. We do not prove it here, but it seems intuitively obvious that firm A is better off in this case than if it prices uniformly. The same is true for firm B (holding firm A's prices as given), by symmetry.
- 4.20 Price discrimination, therefore, is (i) individually optimal but (ii) jointly sub-optimal. Absent coordination or collusion between the two firms, game theory tells us that an agreement not to price discriminate cannot be a competitive equilibrium in this case, as firms would always have an incentive to begin price discriminating and benefit in the short term.

A more general treatment

- 4.21 It is not necessarily the case, however, that best-response asymmetry will always lead to reduced prices, as demonstrated by Corts (1998).
- 4.22 Consider if two firms (A, B) each face two markets (1, 2) with best-response asymmetry: 1 is A's strong market and B's weak market and 2 is A's weak market and B's strong market. Corts (1998) notes that a firm's

²⁹ If consumers are not uniformly distributed but, for example, are more likely to be grouped in the middle, then price discrimination may increase prices for those consumers at the extremes, but decrease prices for those in the middle (see Armstrong (2006), footnote 35). The intuition for this result is straightforward: the middle ground becomes more important for firms, so they compete harder for those consumers in the middle and less hard for those at the extremes.

³⁰ See Stole (2006), page 19, for a summary.


³¹ This raises an interesting issue regarding whether firms can commit to not price discriminating and thus earn higher profits, which we discuss in Section 4.g below.

best-response price³² in its strong market will always exceed its best-response in its weak market:

$$\begin{aligned} BR_1^A(p) &> BR_2^A(p) \\ BR_1^B(p) &< BR_2^B(p) \end{aligned}$$

4.23 In other words, for any given price charged by its competitor B, A will charge a higher price in its strong market (B's weak market) than in its weak market (B's strong market).³³ The effect of this is best shown graphically in the following figure, which is reproduced from Stole (2006).³⁴

Figure 9: Best-response asymmetry



4.24 If the firms are able to price discriminate, the equilibria in each market occur at the intersection of the best-response functions, E1 and E2.³⁵ If firms are unable to price discriminate, then each firm's single best-response function will be somewhere in between its two previous discriminatory best-response functions, depending on the relative weighting of the two markets. That is, if market 1 is much more important

³² A best-response function sets out the price that firm A will optimally set for every price set by firm B.

³³ It is important to remember that, even though we describe markets 1 and 2 as being different markets, the same product or service is provided in each market. The only difference between the two markets is the preferences of the groups of consumers.

³⁴ Stole (2006), page 22.


³⁵ $BR_1^A(p)$ represents the price that firm A would set in market 1 for a given price chosen by firm B in market 1. $BR_1^B(p)$ represents the price that firm B would set for a given price by firm A. Equilibrium occurs where these best-response functions intersect, as at this point neither firm would have an incentive to deviate.

to firm A than market 2, then its uniform best-response function will be closer to $BR_1^A(p)$ than $BR_2^A(p)$.

4.25 Therefore, Corts (1988) demonstrates that for any pair of prices bounded in the parallelogram formed by the four best-response functions in Figure 9 above, there exists a pair of market weightings that supports those two uniform prices as an equilibrium. In other words, any pair of prices in that parallelogram could be the equilibrium outcome in uniform pricing. This means that uniform prices could be lower or higher than (weighted average) discriminatory prices. For example, a pair of uniform prices in the region marked H in the diagram above would result in higher prices in both markets.

4.26 The Hotelling model described above implies that firms care much more about their strong market than their weak market. If so, each firm's uniform best-response function will be closer to its strong market best-response function, as in the figure below.


Figure 10: Prices decreasing with best-response asymmetric price discrimination



4.27 The uniform best-response functions (in red) are close to the strong best-response functions. In this case, the uniform pricing equilibrium involves prices that exceed discriminatory prices in each market, as per Thisse and Vives' (1988) result described above.

4.28 However, discarding the uniform Hotelling model and assuming some other demand function could cause the uniform price best-response function to be closer to each firm's weak best-response function, as in the figure below. This could result in the uniform pricing equilibrium being in the region marked 'L' in Figure 9, meaning that discriminatory prices are higher for all consumers.

Figure 11: Prices increasing with best-response asymmetric price discrimination



4.29 The intuition for this is as follows:

- because firms care a lot about their weak market (for example, because it is larger in volume than the strong market), they keep their uniform price low so as not to price out the weak market, and
- when they can discriminate, they can set a high price for their strong market without worrying about pricing out the weak market.

4.30 Corts (1998) concludes that best-response asymmetry is a necessary condition for all discriminatory prices to be lower than the uniform price, but it is clearly not a sufficient condition. Thus, while as stated by Armstrong (2006) best-response asymmetry is 'the crucial feature that can

cause discrimination to intensify competition',³⁶ it does not follow that best-response asymmetry will always result in more intense competition.


4.31 Whilst there is theoretical ambiguity about the effect on prices, in practice we consider that it is unlikely that best-response asymmetric price discrimination would cause prices to increase. To illustrate the intuition behind this, we consider how the firms weight each market when pricing uniformly. As explained above, Firm A's uniform best-response function will be closer to its discriminatory best-response function for Market 1 than to its best-response function for Market 2, if firm A cares more about Market 1 than Market 2. In economic terms, firm A cares more about Market 1 if the profit opportunity is greater there. The profit opportunity is likely to be a function of two things:

- The strength of the firm in that market (this is the intuition behind the Thisse and Vives (1988) result, where the firm cares most about its strong market, as it is from its strong market that it earns its profits).
- The size of that market (even if firm A is not strong in Market 1, it will care more about Market 1 than Market 2 if Market 1 is much larger).

4.32 Given these criteria, consider the possible uniform equilibria in the figure below, where (i) in the bottom left area marked L the uniform price is below the discriminatory prices and (ii) in the top right area marked H the uniform price is above the discriminatory prices.

³⁶ Armstrong (2006), page 115.

Figure 12: Best-response asymmetric price discrimination



4.33 Applying these criteria:

- If the markets are broadly equal in size, then given both firms care more their strong markets, then the uniform equilibrium is likely to be in H.
- If market 1 is much larger than market 2, then both firms will care more about market 1 (despite market 1 being firm B's weak market) and the uniform equilibrium is likely to be close to E1.
- If market 2 is much larger than market 1, then both firms will care more about market 2 (despite market 2 being firm A's weak market) and the uniform equilibrium is likely to be close to E2.

4.34 We consider that the uniform equilibrium is unlikely to be in L, because it would require both weak markets to be large relative to the strong markets. Given firms rank the markets asymmetrically such that one firm's strong market is the other's weak market, this is not possible.

4.35 We consider, therefore, that the uniform equilibrium would likely be in the top right area marked H in Figure 12 above, meaning that best-response

asymmetric price discrimination is likely to make consumers in aggregate better off, or at the very least not worse off.³⁷

Best-response symmetry in theory

4.36 As with best-response asymmetry, we first discuss a simple Hotelling model, before considering a more general formulation.

Hotelling model with choosiness

4.37 Armstrong (2006) describes the following unit demand Hotelling model where consumers have utility:³⁸

$$U = v - td - p$$

4.38 Two firms are situated at each end of the unit interval $[0, 1]$, along which consumers are uniformly located. A consumer located at x (somewhere on the unit interval) is a distance $d_A = x$ from firm A and $d_B = 1 - x$ from firm B. Distance in a Hotelling model such as this can be interpreted as representing consumers' preferences over two differentiated products or over two differentiated online retailers.

4.39 Suppose a consumer wishes to buy a single unit from either firm A or firm B. If he buys from firm $i = A, B$ his net surplus is:

$$u_i = v - p_i - 0.5d_i$$

4.40 v signifies the underlying value the consumer places on the unit, irrespective of the firm from it is purchased,³⁹ P signifies the price charged by firm i and marginal cost is normalised to zero. x measures the differentiation between A and B. t captures how much utility the consumer loses from having to choose his least preferred option, which Armstrong terms 'choosiness'.⁴⁰

³⁷ We say *likely* as this is clearly based on intuition, rather than being a formal result. Furthermore, whilst it will apply in many cases, it will not necessarily be true in all cases, if demand takes an extreme or unusual form.

³⁸ Armstrong (2006), page 110. Note that this is the same model as discussed above, except that the parameter 0.5 is now replaced by a general parameter t .

³⁹ That is, differentiation between the firms is captured by their location, instead of differences in the underlying value, v .

⁴⁰ The distinction between t and x is not necessarily intuitively straightforward. By way of motivation, consider two customers that are immediately next to each other at location $x = x'$. One consumer has high valuation v' and one consumer has low valuation v'' . t measures how much value the consumer loses from travelling. If the parameter t were not included then (that is, t was just assumed to be 1), then each consumer would lose the same amount of utility from travelling. However, given they have different valuations, one would intuitively expect them to lose differing amounts of value from choosing their least preferred option.

- 4.41 v , t and x vary among consumers. Firms may be able to observe and price discriminate based on each of these, each of which has different implications for prices.
- 4.42 Armstrong (2006) considers price discrimination based on both firms observing t . Each firm would like to charge choosy consumers that are near them more than the uniform price, as the firms know those consumers are unlikely to switch to the other firm. Choosiness is, in this respect, best-response symmetric.
- 4.43 Armstrong (2006) shows that if firms can observe the consumer's choosiness (t) but not location (x , using the notation introduced above), then the equilibrium discriminatory price is t .
- 4.44 The equilibrium uniform price in this case is an average of t .⁴¹ This means that choosy consumers will be charged a higher price and less choosy consumers a lower price, relative to the uniform price, if firms are able to price discriminate.

A more general treatment

- 4.45 The model above considers best-response asymmetry only in the context of Hotelling demand. Here we consider a more general model. The countervailing effects at work with best-response symmetry are best illustrated based on the elasticity faced by the firm, following Holmes (1989). In a symmetric duopoly the elasticity faced by each firm is equal to the sum of the elasticity of the whole market and the cross-price elasticity between the two firms:

$$\varepsilon^{firm}(p) = \varepsilon^{market}(p) + \varepsilon^{cross}(p)$$

- 4.46 Market elasticity is related to the Appropriation Effect, in that it measures the extent to which consumers will not consume the good at all; that is, if the market is inelastic, then firms will be able to appropriate more consumer surplus without consumers taking the outside option of not consuming at all.

⁴¹ More specifically, it is the harmonic mean of t . See Armstrong (2006), page 110.


4.47 Cross-price elasticity is related to the Intensification of Competition Effect, as it measures the extent to which consumers will switch to the rival's product.

4.48 In general, therefore, the prices faced by consumers will depend on the interaction between these two effects, and thus on the relative sizes of the market and cross-price elasticities.

4.49 It is possible, however, to draw a more definitive conclusion, building on the reaction functions discussed above. Consider if both firms (A, B) view market 1 as their strong market and market 2 as their weak market. In this case:

$$\begin{aligned} BR_1^A(p) &> BR_2^A(p) \\ BR_1^B(p) &> BR_2^B(p) \end{aligned}$$

Figure 13: Best-response symmetry



4.50 Note that this diagram differs from Figure 11 above for best-response asymmetry. Here the firms both rank market 1 over market 2, meaning the equilibria are in the top right and bottom left corners.

4.51 The discriminatory equilibria in each market are determined by the intersection of the best-response curves, E_1 and E_2 . The uniform equilibrium could be anywhere in the quadrangle bounded by the four best-response functions, depending on the relative weights the firms place on the markets.

4.52 Regardless of the market weights, however, it is clear that any point in that quadrangle would involve an increase in price for the weak market 2 and a decrease for the strong market 1. Thus, the specific result from the Hotelling model above that discriminatory pricing will benefit some consumers and penalise others is in fact a general result.

Combining best-response symmetry and asymmetry

4.53 Firms may be able to identify multiple characteristics, including best-response symmetric characteristics as well as best-response asymmetric characteristics. Consider, for example, if the firm observes only the consumer's brand preference. As discussed above, this price discrimination based on best-response asymmetric heterogeneity makes all consumers better off. What happens if the firm can now also observe the level of choosiness of the consumers or their switching costs (which are both best-response symmetric)?

4.54 The precise effect depends on, among other things, the form of demand. In general, though, the same conclusions hold as discussed above: the incremental effect of identifying best-response symmetric heterogeneity is to make some consumers better off and some worse off:

- Firms will compete even harder over consumers they know have low switching costs, as they know that these consumers are likely to be tempted by a competing offer. These consumers therefore face lower prices than if switching costs were not observed.
- Firms will compete less hard over consumers they know have high switching costs, as they know these consumers are unlikely to be tempted by a competing offer. These consumers face higher prices.

Best-response symmetry and asymmetry in practice

4.55 It is not straightforward to identify whether a particular type of heterogeneity that is observed is best-response symmetric or asymmetric.

Consider, for example, if firms identify the location of the consumers' home. This could be used to identify consumers from an affluent area, on the basis that they may have a higher WTP. This is clearly best-response symmetric. However, it could alternatively be used to identify those consumers that are close to a competing company's high-street store, which is best-response asymmetric.

4.56 The same is true for observing past-purchase history or web browsing history. This could potentially be used to identify which consumers are sophisticated and likely to search for lower prices, which is best-response symmetric. Alternatively, it could be used to gauge brand preferences, which are best-response asymmetric.

4.57 It is not possible, therefore, to categorise different types of heterogeneity as best-response symmetric or asymmetric. In specific cases, however, it may be possible to observe the firm's behaviour and how it conditions prices, and thus conclude whether the discrimination is best-response symmetric or asymmetric.

Conclusion

4.58 Best-response symmetry is arguably the single most important aspect of any model of price discrimination with competition, in that it determines which of the Appropriation Effect and the Intensification of Competition Effect dominates.

4.59 Whilst it is not possible to categorise different types of heterogeneity as generally being best-response symmetric or asymmetric, in specific cases this may be observable.

4.60 In summary:

- With best-response symmetry, the price will decrease for some consumers and increase for others.
- With best-response asymmetry, the result is theoretically more ambiguous: prices could decrease for all consumers (as is the result for a standard Hotelling model), but depending on the form of demand they could increase for some or all consumers. In practice, however, we consider that best-response asymmetric discrimination would more likely decrease prices.

4. b The heterogeneity observed across consumers

Introduction

4.61 Observed heterogeneity across consumers is the means by which price discrimination is possible. The form and features of this observed heterogeneity can determine the welfare implications of price discrimination, as we demonstrate in the section above on best-response symmetry.

4.62 In this section, we discuss three ways in which observed heterogeneity could vary. We first discuss the features of the heterogeneity that is observed. We then discuss the amount of heterogeneity (that is, what happens when a firm is able to observe an increasing amount of information about a consumer's preferences). Finally, we consider the accuracy of the heterogeneity, by which we mean what happens when the heterogeneity observed is only an imperfect signal of WTP.

Characteristics of heterogeneity

4.63 Based on our review of the literature, we have identified the following three different characteristics of heterogeneity:

- WTP, and in particular the ability of the monopolist to segment the consumers between low and high WTP
- best-response symmetry
- past purchasing behaviour.

Distinguishing between low and high WTP

4.64 Consider if a monopolist is able to observe some heterogeneity that allows it to segment consumers into a high WTP group and a low WTP group. As set out above in Figure 4, the net effect of price discrimination in this case is ambiguous, because the Appropriation Effect and the Output Expansion Effect act in opposite directions.

4.65 However, it is possible to draw general conclusions between the consumer segmentation the firm can achieve and the effect on consumer surplus.

Consider if the firm is able to segment its consumers into those that have very high WTP and those that do not. One would not expect the Output Expansion Effect in this case to be very significant, because the firm needs to know which consumers have low WTP (relative to the uniform price) in order to expand the market. In this case, therefore, consumer surplus would likely decrease.

4.66 Consider if, instead, the monopolist was only able to segment consumers into those that have very low WTP (relative to the uniform price) and those that do not. In this case, the Output Expansion Effect would likely be significant, and consumer surplus would likely increase.

4.67 This intuition can also be shown more formally in a simple model.

4.68 Consider if consumer valuations are uniformly distributed between v_L and v_H , which implies a linear demand curve. Marginal cost is normalised to 0.

4.69 Quantity demanded in this case is equal to:

$$\begin{aligned} &(v_H - p) \text{ if } p \geq v_L \\ &v_H - v_L \text{ if } p < v_L \end{aligned}$$

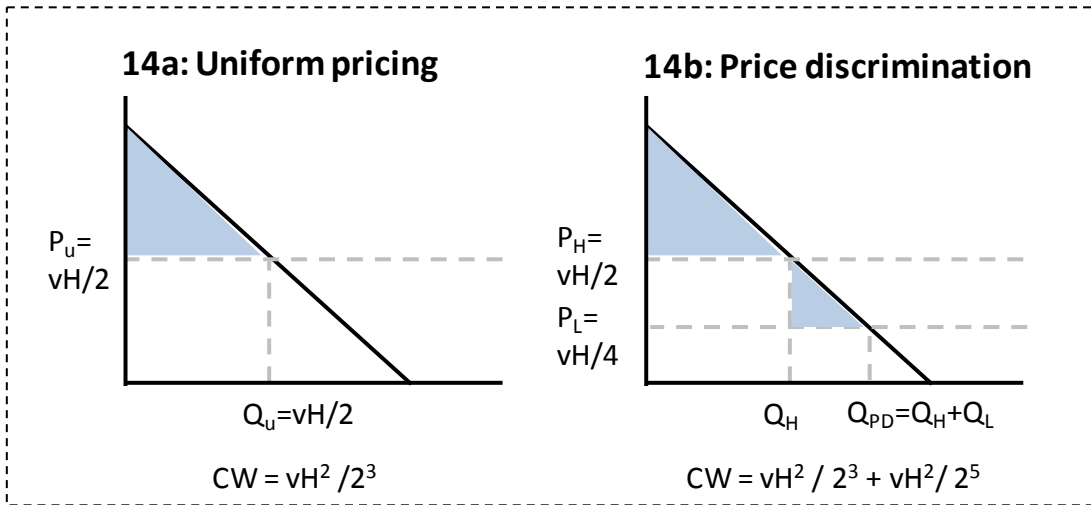
4.70 Optimal uniform prices are therefore:

$$\begin{aligned} p^* &= \frac{v_H}{2} \text{ if } \frac{v_H}{2} \geq v_L \\ p^* &= v_L \text{ if } \frac{v_H}{2} < v_L \end{aligned}$$

4.71 This is the interior and corner optimal price, respectively. The intuition behind it is straightforward: firms will charge the optimal interior price unless that is below v_L , in which case they can increase the price to v_L without losing any customers. v_L therefore acts as a lower bound on the optimal price.

4.72 For tractability, assume that valuations are uniformly distributed between 0 and v_H , meaning that the optimal uniform price is $v_H/2$. Now consider if the monopolist can observe some heterogeneity that allows him to segment consumers into a group with valuation between 0 and $v_H/2$ and a group with valuation between $v_H/2$ and v_H . The optimal prices for the groups are $v_H/2$ and $v_H/4$, respectively. We summarise the effect on consumer surplus in the figure below.

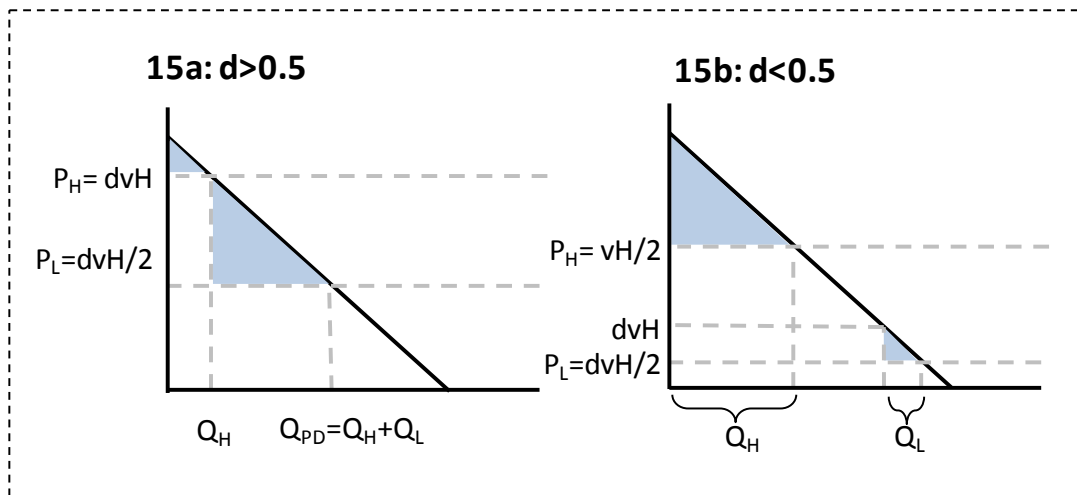
Figure 14: Third-degree monopolist price discrimination



4.73 Note that consumer surplus is increased by the third degree price discrimination (the lower triangle of consumer surplus in Figure 14b is the increase in consumer surplus).

4.74 Consider now a more general model where the heterogeneity observed by the firm allows it to segment the consumers into two groups: a low valuation group with valuation between 0 and dVH and a high valuation group with valuation between dVH and VH , where d is between 0 and 1. We summarise the effect on consumer surplus in the figure below.

Figure 15: Third-degree monopolist price discrimination with different segmenting



4.75 It is possible to relate d and consumer surplus algebraically. If $d < 0.5$:

$$C_W = \frac{v_H}{2^3} + \frac{d^2 v_H}{2^3}$$

$$\Delta C_W = C_{W_{PD}} - C_{W_U} = \frac{d^2 v_H^2}{2^3} > 0$$

4.76 If $d > 0.5$:

$$C_W = \frac{(1-d)^2 v_H^2}{2} + \frac{d^2 v_H^2}{2^3}$$

$$\Delta C_W > 0 \text{ if } d < 0.6, \Delta C_W < 0 \text{ if } d > 0.6$$

4.77 This is a basic model that makes certain simplifying assumptions that may not hold in practice. It does, however, allow us to draw two general conclusions. First, price discrimination of this sort can increase consumer surplus under certain circumstances. Second, the precise effect on consumer surplus depends on how the monopolist can segment the consumers. Consumer surplus decreases if the monopolist is identifying high valuation consumers (d is high), but it is increasing if the monopolist is identifying low valuation consumers (d is low).

Figure 16: Consumer surplus with third-degree price discrimination by a monopolist



4.78 This can also be thought of in terms of the consumers that did or did not purchase under uniform pricing. If a monopolist that prices uniformly is able to identify heterogeneity in existing consumers (who clearly have a relatively high valuation) then the increase in consumer surplus will be lower than if the monopolist is able to identify heterogeneity in potential

customers (whose valuation is not high enough to purchase at the uniform price).

4.79 Note that this conclusion is for a monopolist only. It does not extend, for example, to a duopoly with best-response asymmetry. In this case, identifying a particular subset of consumers will have different implications for each of the firms: that is, one firm is identifying its low WTP consumers whereas the other firm is identifying its high WTP consumers.⁴² It can be shown that discriminatory prices in this case will be less than the uniform price for all consumers, provided that consumers are distributed uniformly.

Best-response symmetry/asymmetry

4.80 We describe above in Section 4.a Armstrong (2006)'s duopoly, unit demand model where consumers have utility:⁴³

$$U = v - td - p$$

4.81 v , t and x (which determines d) vary among consumers (see Section 4.a for a description of what these parameters mean). Firms may be able to observe and price discriminate based on each of these, each of which has different implications for prices.

4.82 First, Armstrong (2006) demonstrates that firms would not optimally price discriminate based on valuation. For simplicity assume that all consumers have the same choosiness parameter, t . Armstrong (2006) shows that each firm will optimally set price equal to t . That is, valuation being observable has no impact on outcomes.

4.83 This demonstrates the importance of the level of competition in the market. As shown above, a monopolist would clearly benefit from knowing the distribution of valuations of the consumer. A duopolist, however, would not take valuations into account even if they were observable. The intuition for this result is that valuation in this model is assumed to be invariant between firms. Thus, whilst it affects the overall level of consumer surplus, it does not affect how firms set prices under competition.

⁴² This result may change if firms cannot identify the same customers, but, for example, can only segment consumers that have purchased from them previously. We discuss this in more detail when we consider dynamic models below.

⁴³ Armstrong (2006), page 110.

4.84 This result is also true when the model is adapted such that consumers buy multiple units and multiple products, and where consumers' preferences over these various units are private information.⁴⁴

4.85 Second, as per Thisse and Vives (1988) discussed above,⁴⁵ best-response asymmetric discrimination based on location (or equivalently brand preference) results in lower prices for all consumers, assuming that consumers are uniformly distributed. As discussed above, however, this is not a general result: with more general demand, best-response asymmetric discrimination could cause prices to decrease or increase.

4.86 Third, Armstrong (2006) shows that best-response symmetric discrimination based on choosiness (t) results in prices decreasing for choosy customers (high t) and increasing for non-choosy customers (low t). As discussed above, this is in fact a general result.

Past purchasing behaviour

4.87 Some contributions to the literature (see Armstrong (2006)) roughly characterise online price discrimination as the ability of firms to observe the past purchasing behaviour of consumers.

4.88 We describe this in detail in Section 4.c below. The conclusions, however, are not significantly different from the general models discussed above:

- If observing past purchase history results in firms identifying best-response asymmetric heterogeneity, then consumer surplus is likely to improve.
- If observing past purchase history results in firms identifying best-response symmetric heterogeneity, then the effect on consumer surplus is more ambiguous.

4.89 As a general rule it is not possible to classify price discrimination based on past purchasing history as being harmful or beneficial to consumers. We discuss the specific models in Section 4.c below.

⁴⁴ Armstrong and Vickers (2011) and Rochet and Stole (2002).

⁴⁵ See paragraph 4.18 above.

Amount of heterogeneity

4.90 The amount of heterogeneity observed by firms is particularly relevant to determining the balance between the Appropriation Effect and the Output Expansion Effect. As above, we first consider a monopolist, before considering an oligopoly.

Monopoly

4.91 In basic terms, as the amount of heterogeneity observed by a monopolist increases, then in the limit the Appropriation Effect becomes increasingly important relative to the Output Expansion Effect, and consumer surplus is likely to decrease. Total surplus will increase as the market expands, meaning that all potential gains to trade are exploited (but then appropriated by the firm).⁴⁶

4.92 This result is intuitively straightforward:

- Under first-degree price discrimination the firm has perfect information about each consumer's WTP, and thus the Appropriation Effect dominates and consumer surplus is eradicated by price discrimination.
- Under third-degree price discrimination, the firm is able to divide consumers into high and low WTP groups based on some observable heterogeneity. The more relevant (that is, informative about WTP) heterogeneity is observed, the more groups the firm can segment. In the limit, fewer and fewer individual consumers are contained in each group, until it approaches perfectly individualised first-degree price discrimination, meaning consumer surplus is reduced.

4.93 The intuition for why this is the case is also straightforward: the Output Expansion Effect is bounded by the number of consumers who have WTP in excess of marginal cost, whereas the Appropriation Effect is not bounded in this way.

4.94 Consider, for example, if a monopolist can already observe significant amounts of relevant heterogeneity. Then, whilst it cannot perfectly price discriminate, it can charge a variety of different groups of consumers

⁴⁶ Varian (1988) shows that price discrimination increases total surplus whenever it causes quantity to increase, as deadweight losses from having prices above marginal cost are decreased.

different prices. Consider the effect of that monopolist being able to observe additional heterogeneity. This is likely to allow the monopolist to appropriate additional welfare. It is unlikely, however, to result in quantity increasing significantly, because low WTP consumers have already been priced into the market by the heterogeneity already observed.

4.95 We present a simple model of this below.


A model of the amount of heterogeneity observed

4.96 Using the model set out above, if $d=0.5$ the monopolist is able to observe heterogeneity that allows it to split all potential consumers into two groups of equal size based on their valuation. In this case, price discrimination increases consumer surplus.

4.97 In this example, the firm is able to observe a single piece of heterogeneity (let's call the number of separate pieces of information θ , so in this example $\theta=1$, whereas with uniform pricing $\theta=0$), such as post code. We are interested in modelling what happens to consumer surplus if the firm is able to observe additional pieces of heterogeneity, so for example if the firm can observe a consumer's post code and whether that consumer has recently referred to a particular product in an email.

4.98 Let us assume that the additional piece of information tells the firm something new about WTP (that is, it does not just divide all possible consumers in the same way as the first piece of information). So, under $\theta=1$ the firm can divide consumers into two equal-sized groups, whereas with $\theta=2$ the firm can divide consumers into four equal-sized groups.


Figure 17: Consumer surplus and amount of information



4.99 If $\theta=3$, then the firm can divide the consumers into eight equal-sized groups, and so on. It is possible to formalise the relationship between θ and consumer surplus:

$$CW = \frac{v_H^2(2^\theta - 1)}{2^{2\theta+1}} + \frac{v_H^2}{2^{2\theta+3}}$$

Figure 18: Consumer surplus and number of observations of heterogeneity



4.100 This is a basic model that makes certain simplifying assumptions that may not hold in practice. Nevertheless, a general conclusions can be drawn that, whilst consumer surplus may increase with some observed

heterogeneity, it then decreases with more observed heterogeneity. As described above, the intuition for this is that the Output Expansion Effect is bounded whereas the Appropriation Effect is not. Second, consumer surplus decreases to zero as θ increases. This is because third-degree price discrimination tends to first-degree price discrimination as θ tends to infinity.

4.101 This is true regardless of the initial value of d . Even if d is such that price discrimination is initially beneficial to consumer surplus (because the monopolist is identifying low WTP consumers), as the amount of information increases consumer surplus decreases to zero in the limit.

Oligopoly

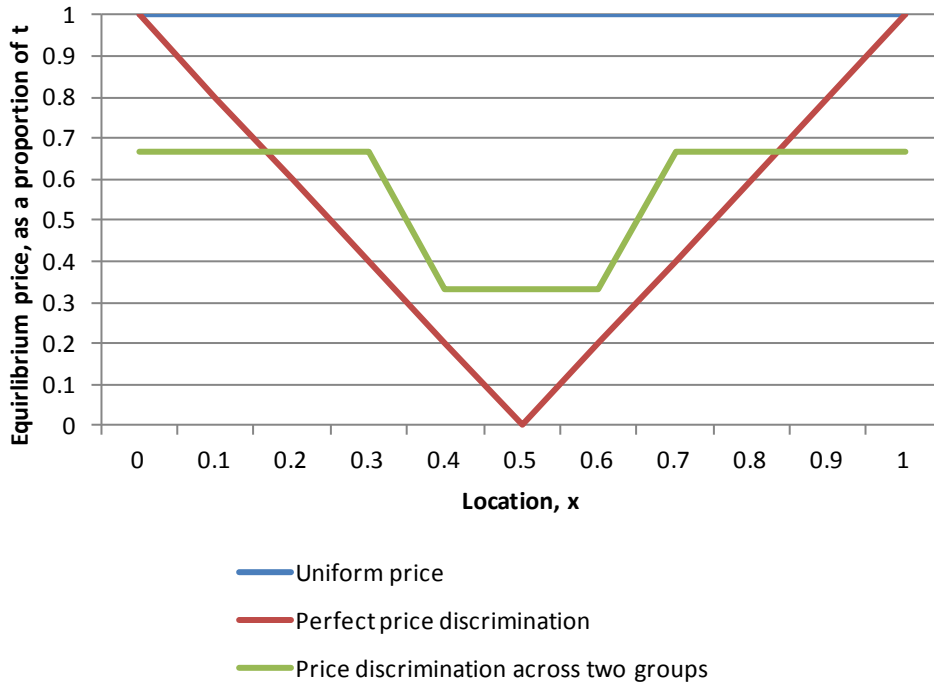
4.102 The Intensified Competition Effect makes the effect of additional heterogeneity in oligopoly models more ambiguous. As described above, there are models where perfect price discrimination causes prices to fall relative to uniform pricing. Intuitively, therefore, it is likely that moving from limited heterogeneity observed to complete heterogeneity observed would cause average prices to fall.

4.103 This can be seen by comparing Bester and Petrakis' (1996) result with that of Thisse and Vives (1988).

4.104 Bester and Petrakis (1996) find that if firms are able to segment consumers into those with a preference for firm A (that is, $x < 0.5$) and those with a preference for firm B ($x > 0.5$) then firm $i=A,B$ will offer a price of $2t/3$ to customers that prefer firm i and $t/3$ to customers that prefer firm $j \neq i$.

4.105 This can be compared to Thisse and Vives (1988), in which firms can directly observe and condition prices on x , allowing them to price discriminate perfectly. The figure below summarises the prices paid by consumers.

Figure 19: Increasing heterogeneity observable to oligopolies



4.106 Consumers face lower prices with price discrimination, in comparison to uniform pricing. The shift from price discrimination across two groups to perfect price discrimination causes some consumers to be worse off (that is, those at the extremes with $x < 1/6$ or $x > 5/6$) and some to be better off (those in the middle, $1/6 < x < 5/6$).

4.107 The intuition for this is as follows: the consumer at, for example, $x=0.1$, benefits from moving from uniform pricing to price discrimination based on two groups, as a result of the Intensification of Competition Effect. He is an infra-marginal consumer, as he is at an extreme. Under perfect price discrimination, he is no longer an infra-marginal consumer, but instead sits within his own market as a marginal consumer.

4.108 Note, however, that the mean price paid is lower with perfect price discrimination than with price discrimination across two groups. It appears, therefore, that increased observable heterogeneity may be beneficial to consumers with best-response asymmetry.

4.109 As described in Section 4.a, the effect of price discrimination with best-response symmetry is, in general, ambiguous, in that it causes discriminatory prices to disperse around the uniform price. The effect of

increased observable heterogeneity is likely to be similarly ambiguous, although we have not seen formal models that consider this issue: the dispersal of the discriminatory price around the uniform price is smaller if firms observe small amounts of symmetric heterogeneity than if they observe large amounts. In either case, the average price paid by the consumer remains unchanged.

Accuracy of observed heterogeneity

4.110 The accuracy with which firms observe heterogeneity and conclude on willingness to pay is also of interest. In some cases, for example, firms may incorrectly conclude that a consumer has high willingness to pay. In this case, the consumer would not purchase the product under price discrimination, but may have done under uniform pricing. As a result, price discrimination causes the consumer surplus and profit associated with that consumer to be lost.

4.111 The effect on aggregate consumer surplus, however, is less clear, as we discuss below. We first discuss a simple monopoly model, before considering oligopoly models.

Monopoly

4.112 Consider if consumers vary in their valuation, such that they can be divided into two groups, L and H, with low and high WTP, respectively. The single firm in the market cannot observe their WTP, but can observe a feature of consumers, γ , which takes one of two values (0 or 1). γ is a noisy indicator of WTP, in that:

$$\begin{aligned} P(\gamma = 1|H) &= p \\ P(\gamma = 0|H) &= 1 - p \\ P(\gamma = 1|L) &= q \\ P(\gamma = 0|L) &= 1 - q \end{aligned}$$

4.113 That is, a consumer with a high valuation has a $p\%$ chance of exhibiting characteristic γ and a $(1-p)\%$ chance of not exhibiting that characteristic. In the two extreme cases:

- **Zero correlation:** If the observable feature is not correlated at all with the WTP of the consumers ($p=q$), then the monopolist has no useful

information on which to condition prices, and therefore just charges the uniform prices to all consumers.

- **Perfect correlation:** If the observable feature is perfectly correlated with WTP ($p=1$, $q=0$), then the firm can accurately price discriminate (note that this does not mean perfect price discrimination, it just means the firm can segment the consumer into two groups perfectly).

4.114 Clearly inaccurate price discrimination falls somewhere in between these two examples. As the accuracy of the signal improves, the result moves from the zero correlation case above towards perfect correlation.

4.115 Note, however, that accuracy in this case has two dimensions:

- **Type 1 error:** incorrectly identifying high WTP consumers as low WTP consumers (using the notation above, this probability is $1-p$).
- **Type 2 error:** incorrectly identifying low WTP consumers as high WTP consumers (using the notation above, this probability is q).

4.116 The effect on consumer surplus on moving from zero correlation to perfect correlation depends on the balance of Type 1 and Type 2 errors. If there are more Type 1 errors than Type 2, then inaccurate price discrimination may be better for consumers than accurate price discrimination or uniform pricing. If there are more Type 2 errors, then the reverse may be true.

4.117 This analysis excludes the responses of firms to inaccurate signals, effectively assuming that the accuracy of the signal is exogenous and not controlled by the firm.

4.118 The key implication of this sketch of a model is that the effect of inaccuracy is not neutral, in that depending on the balance of the inaccuracy between Type 1 and Type 2 errors it may have a direct effect on consumer welfare.

4.119 In practise, however, it is unlikely to be possible to observe the type of error being made, as it would require knowledge of the consumer's true valuation. Firms cannot observe this (otherwise they would condition upon

it), and there is no reason to think that a competition or regulatory authority would be able to.

Oligopoly

4.120 As described in the preceding sections, for some oligopoly models price discrimination is beneficial for consumers. As one would intuitively expect, in these models increasing the accuracy of the price discrimination is beneficial to consumers. Esteves (2012) shows that consumer welfare under imperfect best-response asymmetric price discrimination always exceeds that under uniform pricing, and that consumer welfare is increasing in the level of accuracy with which firms segment consumers. Note that issues regarding Type 1 and Type 2 errors do not apply in this oligopolistic setting, as firms rank consumers asymmetrically.

4.121 Chen et al. (2001) present a similar result and summarise the intuition: best-response asymmetric price discrimination allows firms to compete more intensely for each consumer (as we describe above). If there is a chance that the firm is inaccurately identifying that consumer as preferring the other firm, however, then that softens the price competition for that consumer. This model assumes that the firms have the same level of accuracy.

4.122 We have not found models that consider inaccuracy in best-response symmetric models. Our intuition, however, is that the effect would be analogous to the monopoly situation discussed above, and would depend on the balance of the inaccuracy between misidentifying (i) consumers with high switching costs as having low switching costs and (ii) consumers with low switching costs as having high switching costs.

Conclusion on accuracy

4.123 The effect of errors by the firm in this sense is not straightforward to model. If price discrimination has a positive effect on consumers, such as best-response asymmetric discrimination, then increased accuracy is beneficial for consumers. If price discrimination has an ambiguous effect on consumers (as under monopoly or best-response symmetric discrimination) then what we have termed the balance of the inaccuracy adds another layer of complexity and can affect consumer welfare.

Summary

4.124 For monopoly models, we have shown that:

- Consumer surplus may increase or decrease with price discrimination depending on the balance of the Appropriation Effect and the Output Expansion Effect, but it is more likely to decrease if firms are segmenting high value consumers rather than low value consumers.
- Consumer surplus will decrease if the amount of heterogeneity observed increases such that the monopolist is able to segment consumers into smaller groups, in that third-degree price discrimination tends towards first-degree price discrimination as observable heterogeneity increases.
- The accuracy with which price discrimination is undertaken has an ambiguous effect on consumer surplus.

4.125 For oligopolists:

- Best-response symmetry or asymmetry is the single most important feature of the observed heterogeneity, as discussed in Section 4.a above.
- Increasing the amount of heterogeneity observable to the oligopolists increases aggregate consumer surplus under best-response asymmetry, although some individual consumers may be worse off.
- Consumer surplus increases as the accuracy with which firms price discriminate in best-response asymmetric settings.

4c Dynamic interactions between firms and consumers

Introduction

4.126 Up to this point, we have reviewed static, one-shot models of price discrimination. However, in reality there are many markets in which interactions between consumers and firms are repeated, as consumers buy multiple goods across multiple periods.

4.127 This dynamic introduces new issues into the analysis, such as the extent to which consumers are sophisticated or naïve. Issues that we have discussed already (in particular best-response asymmetry) are also important in these models.

4.128 As in previous sections, we first describe the effects of monopoly dynamic pricing, before considering various oligopoly models.

Monopoly

4.129 Consider a simple dynamic model with (i) a single firm, (ii) N consumers with valuation uniformly distributed between 0 and 1 and unit demand and (iii) two periods, 1 and 2. Possible forms of price discrimination in this context include:

- inter-temporal price discrimination where the price charged in the first period is not the same as the price charged in the second period
- behaviour based price discrimination, where the price charged to a consumer in the second period varies depending on whether the consumer purchased the good in the first period.

4.130 We focus on behaviour based discrimination, rather than inter-temporal discrimination, as inter-temporal discrimination is not a form of personalised pricing, and is thus outside the scope of this review.

4.131 A key assumption within these dynamic models is the sophistication of the consumer. We first discuss naïve consumers, before considering sophisticated consumers.

Naïve consumers

4.132 The monopolist observes whether the consumer purchased in the first period. It does not observe any other information about the consumer. Consider if consumers are naïve, and do not take into account the possibility that their purchasing decision in period 1 affects the price they are offered in period 2.

4.133 In this case, a consumer will purchase in the first period if the price he receives (p_1) is less than his valuation. This then reveals to the monopolist

that $v > p_1$, meaning that the monopolist will charge a higher price to that consumer in the second period. Using the terminology introduced earlier, by purchasing in the first period the consumer reveals himself to be in the monopolist's strong group. Similarly, consumers that do not purchase in the first period reveal themselves to be in the monopolist's weak group.

4.134 In this situation, Armstrong (2006) shows that consumer surplus will be lower than if the firm cannot price discriminate. This is because the naïve strong consumer allows himself to be exploited by the monopolist in the second period, when he is charged a higher price as a returning customer. There is still an Output Expansion Effect through weak consumers (those that did not purchase in the first period) being priced into the market, but this is dominated by the appropriation of consumer surplus from the naïve strong consumer.

Sophisticated consumers

4.135 If the consumer is sophisticated, however, he realises that by refusing to purchase in the first period (even if $p_1 < v$) he could get a lower price in the second period. In this case, the firm's incentive to price discriminate depends on the extent to which the monopolist can commit to its second period price in period 1.

4.136 If it cannot commit to its second period price, then the sophisticated consumer knows that the monopolist will always have an incentive to charge a lower price to non-returning customers in the second period. The monopolist in this case sets a relatively low first period price in order to prevent high value consumers simply waiting and being incorrectly identified as low value consumers in the second period. The sophisticated consumer is essentially able to extract a low first period price, making him better off and the monopolist worse off.

4.137 If, on the other hand, the monopolist is able to commit to prices, then it commits not to price discriminate, given that it is worse off under price discrimination. This comes from the Commitment Effect discussed in Section 3, and is analogous to the Coase conjecture regarding dynamic monopoly pricing (Coase, 1972).

Summary

4.138 The effect on consumer surplus in dynamic monopoly models therefore depends on (i) consumer sophistication and (ii) commitment ability.

Table 3: Summary of effects in dynamic games

	Monopolist able to commit to prices	Monopolist unable to commit to prices
Sophisticated consumers	Monopolist chooses not to price discriminate	Price discrimination increases consumer surplus
Naïve consumers	Price discrimination harms consumer surplus	Price discrimination harms consumer surplus

4.139 This is based on Armstrong (2006), but Acquisti and Varian (2001) present a similar model with similar results. In particular, they find that the profitability of price conditioning across time periods depends on the proportions of consumers that are sophisticated and naïve. If a sufficiently high proportion of consumers are naïve, then price conditioning in this way would be profitable.

4.140 The sophistication of consumers is clearly an important issue. We discuss this in detail in Section 4.f below, including an assessment of whether consumers are likely to be sophisticated or naïve in practice. The general indication appears to be that consumers are not particularly sophisticated (in the sense in which these dynamic models define it).

Oligopoly

4.141 There are two strands of the literature that examine price discrimination based on past purchasing history:


- models where purchase history reveals information about a consumer's **brand preferences**, which is best-response asymmetric
- models where purchase history reveals information about the **switching costs** consumers face, which in the case we consider is **best-response symmetric**
- more complex models of oligopoly dynamic pricing.

4.142 The distinction between models of brand preference and models of switching costs can be thought of in terms of differentiation. Models of brand preference assume that firms are ex-ante differentiated, whereas models of switching costs assume that firms are ex-ante homogenous but become differentiated ex-post due to switching costs (this could be applied, for example, to credit cards or telephone services).

Brand preference

4.143 Fudenberg and Tirole (2005) present a model in which the Bester and Petrakis (1996) Hotelling model discussed in Section 4.b above is repeated over two periods.

Figure 20: Hotelling model



4.144 Without going into the technical details again, the Bester and Petrakis (1996) model involved best-response asymmetric price discrimination triggering increased competition and lower prices for all consumers.

4.145 In the Fudenberg and Tirole (2005) model, the two firms observe whether the consumer previously purchased from firm A or firm B in period 1. This tells them whether the consumer has a preference for brand A or brand B, based on which they then price discriminate in the second period.

4.146 The effect on prices depends, as in the example above, on the sophistication of consumers. Sophisticated consumers will realise that in the second period they will be discriminated against based on their purchase in the first period. That is, they realise that if they accept a price from their least preferred firm (the firm furthest away from them) in the first period, then in the second period their preferred firm will incorrectly identify them as being in their weak group and offer them a low price. Consumers are therefore less sensitive to the price charged by their most preferred firm (relative to the static one-shot game), meaning that the price they face in the first period will be higher. In the second period, they can then benefit

by switching to the other firm,⁴⁷ and taking advantage of the lower price they would be charged. Consumers are better off as a result of price discrimination.

4.147 This is a variation of the Intensified Competition Effect: consumers know that in the second period they will receive a poaching price.

4.148 Consider if the consumers are instead naïve. Naïve consumers, unlike sophisticated consumers, are not less sensitive to first period price rises, as they do not consider that they can buy from their least preferred provider and then benefit from lower prices from their preferred provider in the second period. They therefore receive a lower price than sophisticated consumers in the first period. In the second and final period, naïve and sophisticated consumers receive the same price (as clearly there is no distinction between the two in a single period).

4.149 Interestingly, therefore, naïve consumers are actually even better off than sophisticated consumers. The intuition for this is that sophisticated consumers accept a higher price in the first period because they know that they will be better off in the second. Naïve consumers do not consider this, and so do not accept a lower price. Nevertheless, they are still better off in the second period, like sophisticated consumers.

4.150 Crucially, firms understand that naïve consumers will not accept a lower price, and tailor their first-period price appropriately. An important condition for naïve consumers to be better off than sophisticated consumers is that the firm changes their optimal behaviour in some way when dealing with naïve consumers. If the firm does not change its behaviour for naïve and sophisticated consumers then sophisticated consumers will generally be better off than naïve consumers.

4.151 Villas-Boas (1999) considers a similar model, but in an infinite-period overlapping generations context. He finds that price discrimination causes prices in all periods to decrease (whereas under Fudenberg and Tirole (2005) discriminatory prices are higher than uniform prices in the first period and lower in the second). Consumer surplus is improved by price discrimination in both models.

⁴⁷ This means that consumers in the middle of the Hotelling line switch firms between periods 1 and 2. This is inefficient switching, from a total surplus perspective, as clearly at some stage they will be buying from the firm that is further away from them.

Switching costs

- 4.152 Consider if there are two firms and N consumers looking to purchase one homogenous unit of a good in each of two periods. If consumers change supplier between period 1 and 2, they incur a switching cost s , which is uniformly distributed between 0 and t . This model is due to Chen (1997), and can be extended to multiple periods with similar results (as per Taylor (2003)).
- 4.153 In the second period, firms can charge a high price to existing customers and earn a positive profit, as they know that the consumer incurs switching costs (although they do not know exactly what s is for each consumer). Some consumers will switch, particularly if they have a low s or are relatively indifferent between firms.
- 4.154 In the first period, firms will compete to gain market share, which allows them to earn a higher profit in the second period.
- 4.155 The intensity of competition in both periods is determined in part by whether firms can discriminate based on observing first period purchases. If firms can price discriminate in this way, then they can identify consumers that purchased from the other firm and charge them a lower price, on the basis that they are likely to be in their weak group. The ability to price discriminate triggers more intense competition in an analogous way to the static model of Thisse and Vives (1988): firms can make attractive offers to their weak group of consumers without having to decrease the price they offer their strong group.
- 4.156 Firms are clearly worse off as a result of being able to price discriminate. The effect of price discrimination on consumers is actually ambiguous, because of the presence of the switching costs: without price discrimination, consumers would not switch between the first and second periods (as obviously the prices would not change between the first and second periods, given that the second period is just an unrelated repetition of the first), and thus avoid incurring switching costs. Under price discrimination, prices are lower in the first period, but there is additional switching (with the associated costs to consumers) in the second period.

Other models

- 4.157 The two models discussed above suggest that consumers may not be worse off under history-based price discrimination, and may in fact be better off (as noted by Chen (2005), in his general review of history based price discrimination) due to the Intensified Competition Effect. This leads Gehrig, Shy and Stenbacka (2012) to conclude that 'a policy of banning history based price discrimination leads to higher prices in most consumer segments, and tends to reduce overall consumer surplus'.
- 4.158 There are, however, exceptions to these models that indicate that these conclusions are not entirely general.
- 4.159 Consider, for example, if firms have imperfect information about each other's customers: that is, firm A can identify its own strongest customers based on their purchase history, but cannot identify firm B's strongest customers (Shin and Sudhir (2010)). This will be the case where firms observe the purchase history at their own firm, but not at other firms (which intuitively is likely to be the case in many situations).
- 4.160 In this case, price discrimination does not lead to an Intensified Competition Effect, as firms are simply segmenting their strong group of consumers (those that have already purchased from them) into two sub-groups: a very strong group and a strong group. In this case, price discrimination reduces consumer surplus.
- 4.161 The literature also considers price discrimination through long term contracts. We do not consider this here as it does not seem to be applicable to the types of online purchases on which we focus. See Chen (2005) for a summary of these types of models.

Summary

- 4.162 The general static results discussed in the preceding sections broadly apply to dynamic models as well. In particular, where firms identify previous purchasing behaviour as an indicator of brand preference, then this kind of best-response asymmetric discrimination is usually beneficial for consumer surplus.

4.163 Dynamic models introduce additional detail regarding the extent to which consumers are sophisticated or naïve. In most monopoly models, sophisticated consumers are likely to be better off than naïve consumers. If there is competition then naïve consumers are protected to a certain degree, such that in some models they are actually better off than sophisticated consumers. We consider naïve and sophisticated consumers in more detail in Section 4.f.

4.d Personalised advertising

Introduction

4.164 There is significant literature regarding personalised advertising. Hoffman, Inderst and Ottaviani (2012) draw a distinction between targeted advertising, which allows firms to restrict advertising to customers that are predisposed to buy their product, and tailored advertising, which results in firms revealing different product attributes to different consumers.

4.165 We do not discuss personalised advertising in detail, as it is not within the scope of this report and has, in any case, been the subject of other OFT reports.⁴⁸ We discuss it only in so far as it directly affects price discrimination. We discuss targeted and tailored advertising in turn.

Targeted advertising

4.166 The dynamic models discussed above (including, for example, Chen (1997) and Fudenberg and Tirole (2000)) assume the consumer has perfect information. If this is not the case, then the firm can reveal certain information to consumers. Mass advertising involves revealing the same information to all consumers. Targeted advertising involves revealing different information to different subsets of consumers.

4.167 There is significant evidence that firms do currently target their advertising based on data collected about consumers: OFT (2010), for example, estimates that this form of advertising was worth between £64 million and £95 million a year.

⁴⁸ OFT (2009) 'Advertising of Prices' and OFT (2010) 'Online Targeting of Advertising and Prices'.

Direct advertising

4.168 Iyer, Soberman and Villas-Boas (2005) consider a model in which some consumers care about branding and some (which the authors term 'shoppers') do not care about branding at all and only care about price. If firms are able to price discriminate between the two groups, then they will charge a higher price to those that care about branding and a lower price to those that do not (this is because the heterogeneity is, using the terminology introduced in Section 4.a, best-response symmetric). Price discrimination therefore has no effect on aggregate prices or profits.

4.169 If firms can target their advertising as well, then this result no longer holds, and prices are higher than under uniform advertising and pricing. The intuition for this is that targeted advertising allows firms to increase the level of differentiation in the market by sending fewer messages to shoppers, who therefore have fewer opportunities to compare prices.

4.170 Interestingly, this result is independent of whether firms can price discriminate or not. That is, under Iyer et al's (2005) model, gaining the ability to price discriminate when the firm already has the ability to target advertising has no effect.

Platform

4.171 De Cornière and de Nijs (2013) set out a model in which platforms (such as Google or Facebook) sell advertising space to firms. The platforms have access to information about heterogeneous consumers, and can choose to disclose this information to the firms.

4.172 If platforms do not disclose the information, then firms will bid for advertising slots, and those slots will be assigned randomly to consumers. There is no personalisation or targeting of the advert seen by any individual consumer.

4.173 If platforms disclose the information, then firms submit higher bids for access to the consumer(s) with the highest valuation for their products. Consumers therefore do not view irrelevant adverts, but only adverts for products in which they are interested. That is, the use of targeted advertising improves the matching process between firms and consumer.

The firm also charges a higher price for the product he is advertising, as he knows that the consumer seeing the advert is in his strong market.

- 4.174 The overall effect on consumer surplus is ambiguous, as consumers (i) cease to see irrelevant adverts, and thus are matched with their most preferred products and (ii) pay a higher price for the product they purchase. That is, they receive a higher quality product (or at least a product they prefer more), but they also pay more for it.
- 4.175 One implication of de Cornière and de Nijs (2013) is that, as targeted advertising becomes more prevalent, firms expect to reach more consumers in the strong markets, and fewer from their weak markets. The uniform price that they charge thus moves closer to the price they would charge to their strong market consumers if they could price-discriminate. When the number of potential advertisers becomes large, the uniform price equals the strong market discriminating price.
- 4.176 In this case, the firm can effectively achieve the same outcome as under price discrimination simply by targeted advertising: it identifies high value consumers and charges them a higher price. Thus the ability to condition prices directly on observable heterogeneity provides little incremental benefit over targeted advertising.

Summary

- 4.177 Under both the models discussed targeted advertising results in the same outcome as price discrimination without actually conditioning prices on observed heterogeneity, meaning that price discrimination makes little to no difference. Interestingly, this may suggest that price discrimination may not be as common as may be expected.

Tailored advertising

- 4.178 Hoffman, Inderst and Ottaviani (2012) present a model in which firms are unable to communicate all the attributes of a particular product, due to restrictions in airtime, advertising space or consumers' limited attention. Instead, the firm communicates a subset of the product attributes.
- 4.179 If firms are able to observe the particular preferences of consumers, they will then choose which attributes to optimally communicate to individual consumers. If firms are not able to observe any heterogeneity, then they

will non-selectively communicate the same subset of attributes to each consumer.

- 4.180 The welfare implications of this model depend on what is assumed regarding:
- the sophistication of the consumers (and in particular whether they understand the firm's strategy in revealing certain attributes and not revealing others)
 - the ability of the firm to price discriminate, as well as tailor advertising
 - the level of competition in the market.
- 4.181 We illustrate the interaction with price discrimination, by comparing the effect of tailored advertising on consumer surplus with and without the ability to price discriminate.

Firms are unable to price discriminate

- 4.182 Hoffman, Inderst and Ottaviani (2012) conclude that if firms can tailor their advertising but cannot price discriminate (possibly because of resale arbitrage), then tailored advertising is likely to increase consumer surplus.
- 4.183 This is unambiguously the case for sophisticated (or 'wary' in the authors' terminology) consumers that take into account the fact that firms will have selectively revealed attributes to them. This is because targeted advertising gives consumers information about a particular attribute, but sophisticated consumers are also able to infer information from what the firm chooses not to reveal to them. Additional information in this case makes the sophisticated consumer better off.
- 4.184 The perceived valuation of naïve (non-wary) consumers (that is, consumers that do not anticipate that firms will selectively tailor the product attributes they communicate) is inflated by tailored advertising, because the consumers do not anticipate that the firm will have revealed the attribute(s) that is most attractive to them and hidden the least attractive. However, if this is true for all firms in the market, then it does not affect how the naïve consumer chooses between firms. In this case, competition between firms means that the absolute over-inflation of

valuation does not matter, as only relative differences between firms matter under competition.⁴⁹

Firms are able to price discriminate

4.185 If firms are able to use the information they gather to price discriminate as well as tailor advertising, then consumer surplus is likely to be harmed for both sophisticated and naïve consumers.

4.186 The perceived valuation of naïve consumers (that is, consumers that do not anticipate that firms will selectively tailor the product attributes they communicate) is inflated by tailored advertising, because the consumers do not anticipate that the firm will have revealed the attribute(s) that is most attractive to them and hidden the least attractive. This inflated value is then appropriated by the firms using price discrimination.

4.187 If consumers are sophisticated, then tailored advertising sophisticated consumers to perceive increased differentiation between the firms, which effectively dampens competition (similar to the model of Iyer et al (2005) discussed above). The perceived differentiation varies by consumer (it is, in that sense, a form of customer heterogeneity that has been induced by the firm), and thus can be used by the firm to appropriate surplus through price discrimination.

4.188 Interestingly, in this case, the authors show that sophisticated consumers are worse off than naïve consumers, because tailored advertising does not increase perceived differentiation for naïve consumers.

Summary

4.189 Personalised advertising and price discrimination are clearly closely linked, in part because they require the same types of information.

4.190 It is difficult to draw general conclusions regarding the interaction between targeted advertising and price discrimination, and in fact there are relatively few models that seek to do so. One recurring result is that targeted advertising can increase the perceived differentiation of firms,

⁴⁹ To develop the intuition for this, see paragraph 0 above, in which we note Armstrong's (2006) result that competing firms have no incentive to discriminate based on the consumer's (firm-invariant) valuation, as this is the same for all firms.

softening competition and possibly decreasing consumer surplus. The ability to price discriminate may or may not exacerbate this, depending on the assumptions of the model.

4.191 However, this fails to consider the various benefits of targeted advertising, including reduced wastage through delivering advertising to uninterested consumers, which would act to increase consumer surplus.

4.192 Interestingly, there are models in which targeted advertising results in the same outcome as price discrimination, without firms actually needing to condition prices on observable heterogeneity.

4.193 In general, however, we do not find a relationship between targeted advertising and the effect of price discrimination on consumer surplus.

4.e Costs associated with price discrimination

Introduction

4.194 There does not seem to be extensive material in the literature that considers the costs associated with price discrimination. In all the models discussed so far, firms have been able to price discriminate costlessly. Furthermore, the models do not consider the ability of consumers to avoid price discrimination at some cost to themselves.

4.195 These are clearly unrealistic assumptions, as collecting (or purchasing) and processing consumer information is likely to involve a cost to the firms. Equally, consumers may be able to avoid price discrimination at a time or monetary cost (there exists, for example, software that prevents information being collected through cookies).

4.196 We briefly consider the cost to firms and consumers in turn.

Cost to firms

4.197 As in previous sections, we first consider monopoly pricing, before considering oligopoly pricing.

Monopoly

4.198 The effect of price discrimination in static monopoly models depends on the balance between (i) the Output Expansion Effect and (ii) the Appropriation Effect. Price discrimination could increase consumer surplus if it allows the firm to profitably sell to low value consumers.

4.199 Consider, however, if obtaining the ability to price discriminate resulted in an increase in the marginal cost faced by the firm. This would limit the Output Expansion Effect (relative to the case with no increase in costs), as even if they are able to price discriminate firms will not sell to consumers whose valuation is below marginal cost.

4.200 There will, therefore, be situations in which firms have an incentive to incur additional costs to price discriminate because it gives them the ability to appropriate additional consumer surplus (even accounting for the fact that they incur additional costs) from high value consumers, without there being a corresponding increase in output. In that sense, costly price discrimination is worse for consumer surplus than costless consumer surplus.

4.201 This analysis comes with two caveats.

4.202 First, costs associated with price discrimination may be fixed rather than variable (such as investing in the IT equipment and expertise needed to acquire cookies) rather than a variable cost. Fixed costs generally do not impact on the pricing of a firm. However, there are circumstances in which one might expect the cost to be variable, such as if the firm were buying customer data on an ongoing basis from a data broker.

4.203 Second, as discussed above, firms may also collect and process data to tailor their advertising, meaning that whilst consumer surplus from pricing may decrease, consumers may benefit from targeted or tailored advertising, as discussed in Section 4.d above.

Oligopoly

4.204 The effect of the cost on consumer surplus is analogous to the discussion of monopoly pricing above: increased costs are likely to result in increased prices.

4.205 Consider, for example, the Thisse and Vives (1988) model discussed in Section 4.a above, in which best-response asymmetry results in all consumers facing lower prices. This occurs because the strong firm for a particular customer (say this is firm A), sets its price such that the weak firm (B) cannot profitably serve that customer. In the example discussed above, in which cost was normalised to zero, this meant that P_A was set such that the consumer would prefer A even if $P_B=0$. If price discrimination increased marginal cost from 0 to c , then clearly firm A would only have to reduce P_A to the level where $P_B=c$, as firm B would never price below that level. In this case, all consumers would face a higher price than in the costless price discrimination model.

4.206 Costly price discrimination of this kind may still be beneficial to consumers, unless the increase in cost is sufficiently high that the Intensified Competition Effect is cancelled out.

Costs to consumers

4.207 Consider if a consumer is able to incur a monetary or time cost c to avoid price discrimination. 'Avoiding' price discrimination in this context could mean various things, such as (i) presenting a blank profile to the firm or (ii) presenting a profile that is similar to the average consumer or even a low WTP consumer.

4.208 We have not found extensive coverage of this issue in the economic literature, and there are no definitive conclusions as to whether the ability of consumers to avoid price discrimination is beneficial or harmful to consumer surplus. We discuss two specific issues that help highlight this ambiguity: the first explains how avoidance may act as a signal to the firm, and the second considers how avoidance is used if consumers are made better off by price discrimination.

Avoidance as a signal

4.209 The firm anticipates that consumers may be able to anonymise themselves in this way, and sets its prices accordingly. It does not follow, therefore, that anonymisation necessarily results in the consumer being charged a low price. Consider, for example, if the consumer incurred c in order to present a blank profile. The firm knows that this could be because the consumer is a new consumer, or it could be because the consumer

has anonymised itself. It will set an optimal price for that group, which will not necessarily be a low price.

4.210 In fact, given that anonymisation incurs a cost to consumers of c , the fact that the consumer has anonymised himself is an indication that his WTP is sufficiently high for it to be worthwhile for him to incur c . Anonymisation in this sense could be a signal of high WTP, and thus result in higher prices being faced by those with blank profiles.

4.211 Similarly, if a firm knows that consumers presenting a low WTP profile could be genuinely low WTP consumers or high WTP consumers that have anonymised themselves then it will charge a higher price than if it were only genuinely low WTP consumers.

Avoidance where price discrimination improves consumer surplus

4.212 Consumers are likely to be better off in oligopoly models with best-response asymmetric price discrimination, as discussed above. For example, as discussed above, price discrimination based on purchase history benefits consumers purchasing from a monopolist in a dynamic context.

4.213 Consider, however, an individual consumer that can costlessly delete his purchase history. In that case, the consumer would optimally purchase in the first period and then in the second period (in which he is mistaken for someone who did not purchase in the first period, and so is offered a low price).

4.214 If every consumer did this, then clearly price discrimination would break down and the monopolist could credibly commit to uniform pricing. Consumers would then be worse off.

4.215 The ability to anonymise themselves therefore results in a prisoner's dilemma that makes all consumers worse off: holding the behaviour of other consumers fixed, it is optimal for an individual consumer to avoid price discrimination, which results in an equilibrium in which everyone avoids it.

4.216 Conitzer, Taylor and Wagman (2012) consider such a model, and find that increasing the cost of anonymity can benefit consumers, for the reasons discussed above.

4.217 Conitzer, Taylor and Wagman (2012) also note that the cost of anonymisation may be controlled (at least partially) by the firm: if the firm employs very sophisticated discriminatory techniques, then avoidance is likely to be more costly. In this case, the authors conclude that the firm will set the cost at its optimal level, which is likely to make consumers worse off.

Summary

4.218 Where firms have to incur costs to price discriminate, then that will decrease consumer surplus relative to a benchmark of costless price discrimination. This could cause price discrimination to have a net negative effect on consumer surplus, depending on the precise circumstances involved. Nevertheless, incurring significant costs to price discriminate is an example of a factor that, all other things being equal, makes it more likely that consumer surplus is harmed.

4.219 The cost to consumers of anonymisation is more complicated. The availability of free or cheap anonymisation is not, as one might intuitively expect, necessarily beneficial to consumers. In general, however, there is insufficient evidence to draw concrete conclusions regarding the effect of anonymisation costs.

4.f Consumer sophistication

Introduction

4.220 Consumer sophistication is an important consideration in many of the models that have been discussed in previous sections. It clearly has significant implications for the effect of price discrimination on consumer surplus, but it also has implications for the remedies that could be imposed. For example, it may be that naïve consumers are worse off under price discrimination. This also means, however, that a remedy of releasing additional information of some kind (such as requiring firms to notify consumers that they discriminate) may have no effect if naïve consumers fail to take this into account in their decisions.

4.221 In this section, we first discuss consumer sophistication in theory, bringing together the theoretical implications discussed so far. We then discuss consumer sophistication in practice, and attempt to reconcile the theory with how consumers actually behave in an online market. Finally, we discuss how sophistication relates to the transparency of price discrimination online.

Consumer sophistication in theory

4.222 Consumer sophistication is a general term that covers a broad range of requirements on the part of the consumer, ranging from (i) simply understanding that price discrimination is occurring to (ii) understanding all the retailers' strategies and incentives, and optimising accordingly. We illustrate this distinction by going through the three classes of models discussed above: (i) static models, (ii) dynamic models and (iii) models including advertising.

In static models

4.223 We have not found models in the literature that explicitly link third-degree price discrimination and consumer sophistication in a static context. That is not to say, however, that sophistication is irrelevant in these settings.

4.224 Consider, for example, Salop and Stiglitz's (1977) model of sophisticated consumers that are aware of the distribution of prices between stores and naïve consumers that are not and just choose a store at random. Salop and Stiglitz (1977) show that the naïve consumer will generally end up paying a higher price.

4.225 Varian (1980) examines a similar model with sophisticated and naïve consumers that broadly follow Salop and Stiglitz's (1977) definition, but in an intertemporal context. He finds that stores may randomise their prices in order to second degree price discriminate between sophisticated and naïve consumers. The intuition for this is that the sophisticated consumer will only shop when prices are low (that is, during sales) whereas the naïve consumer will shop anytime. The naïve consumer, as under the Salop and Stiglitz (1977) model, pays a higher price.

4.226 While these models do not explicitly consider third-degree price discrimination, these features could equally apply to static models of price

discrimination. Consider, for example, if in the Thisse and Vives (1988) model consumers were naïve or sophisticated independent of their valuation, brand preference or choosiness. If naïve consumers simply chose randomly between the stores (and do not observe the price charged by the other store), then clearly the Intensified Competition Effect will not operate fully.

In dynamic models

4.227 In the dynamic models we discuss above, consumer sophistication requires more than simply an awareness that price discrimination is occurring. In these models it does not suffice that the consumer knows that price discrimination occurs, and therefore shops around. Sophistication in this case requires that consumers know (i) that firms are observing whether or not they are prior customers and (ii) how their behaviour in the first period affects the price they are charged in the future.

4.228 This is clearly a much more onerous requirement than simply being aware that price discrimination exists.

4.229 By and large, in dynamic models sophisticated consumers tend to be better off than naïve consumers (such as Armstrong (2006)), although there are exceptions (such as the Fudenberg and Tirole (2005) model discussed above, in which naïve consumers are more price sensitive because they fail to realise how they may be better off in the second period).

In advertising models

4.230 The advertising models described above require an even higher level of sophistication: in the Hoffman, Inderst and Ottaviani (2012) model, sophisticated consumers (i) understand that firms selectively reveal certain attributes to them and (ii) are able to infer information about what is not revealed to them.

4.231 If firms are not able to price discriminate, then sophisticated consumers are better off. If firms are able to price discriminate, then this is one of the class of models in which naïve consumers are actually better off than sophisticated consumers (because their naiveté protects them from firms introducing differentiation between firms through advertising).

Consumer sophistication in practice

4.232 Consumer sophistication is clearly important in theory, as the welfare implications can be very different for naïve and sophisticated consumers. Deciding whether consumers are in principle likely to meet the theoretical criteria for being sophisticated is, however, difficult.

4.233 While we have not found experimental or empirical evidence that relates specifically to price discrimination, there is a large body of literature that assesses how consumers draw inferences about unobserved variables (which is what sophistication generally means in the context of the models discussed in this report).

4.234 In general, consumers are not particularly sophisticated when it comes to their beliefs or the inferences they draw:

- Weizsacker (2010) presents a meta-analysis of previously undertaken experiments, and shows that consumers place excessively high importance on their own private observations relative to the publically observable choices of others, meaning that the assumptions of rational expectations are violated.
- Charness and Levin (2004) run a laboratory experiment to test how people update their beliefs when they receive new information, and find that around 50 per cent of all decisions taken were inconsistent with Bayesian updating (which is how one would expect the rational, sophisticated consumer to behave).
- OFT (2010) shows that people exhibit biases in response to how prices are framed. For example, if prices are 'dripped' (meaning the true final price is only revealed at the end of the shopping process once surcharges have been added to the base price), then consumers do not make optimal decisions.

4.235 As well as sophistication of beliefs, preferences are also important. In all the models described above it is assumed that consumers optimise so as to maximise their expected surplus, measured in monetary terms. In reality, consumers are likely to have a broad range of behavioural 'biases', in that they care about other things such as fairness (see OFT (2011) for a summary of relevant behavioural biases in competition).

4.236 Note that, as described above, the definition of what constitutes sophistication varies between models, meaning that whether or not real-world consumers are likely to meet this benchmark may also vary. Despite the biases described above, many consumers are likely to realise that firms may observe features about them and condition their prices accordingly, for example.

4.237 It does not follow, however, that consumers can fully optimise their online behaviour over multiple periods and multiple online purchases, which is the requirement for sophistication in Armstrong's (2006) model.

4.238 It is difficult, therefore, to draw definitive conclusions regarding whether consumers are naïve or sophisticated for the purposes of these price discrimination models. The complexity of the interaction and the likely ability of the consumer to deal with that complexity could be assessed on a case by case basis.

Consumer sophistication and transparency

4.239 We note also that there is a link between the analysis of sophistication and of the transparency of price discrimination: a sophisticated consumer who is unaware that price discrimination is occurring is analogous to a naïve consumer. To put it another way, some of the models require the consumers to take some action (or at least be expected to take some action by the firm) for price discrimination to have positive effects. This action could be to check prices at other stores or to reduce their expenditure to pretend to be a low value consumer. Clearly, consumers might not take these actions if they are unaware that price discrimination is occurring.

4.240 This is not explicitly addressed in the models we have reviewed, primarily because most equilibrium solution concepts assume that consumers' expectations are correct in equilibrium. Nevertheless, as a rule of thumb, price discrimination that is not transparent is unlikely to be better for consumers than price discrimination that is transparent.

Summary

4.241 As a general rule, naïve consumers are likely to be worse off than sophisticated consumers, for the simple reason that the sophisticated

consumer always has the option of imitating a naïve consumer (although we do discuss models in which this is not possible, and the naïve consumer is better off).

4.242 The only exceptions to this rule occur when the firm changes its behaviour when selling to naïve consumers (as in Fudenberg and Tirole (2005), for example). Clearly, this requires that the firm is able to segment naïve and sophisticated consumers.⁵⁰ If there is a mixed population of sophisticates and naïfs and the firm cannot segment them, then sophisticated consumers will always be better off than naïve consumers.

4.243 It does not follow, however, that naïve consumers are therefore worse off as a result of price discrimination. There are models in which naïve consumers are better off due to the Intensified Competition Effect, just not as well off as sophisticated consumers.

4.244 Nevertheless, all other things being equal, consumers may be more likely to be harmed by price discrimination if they are naïve. It is not possible to draw definitive conclusions regarding whether consumers are, as a rule, sophisticated or naïve (although most results from the experimental literature suggest consumers are not particularly sophisticated). Instead, this would need to be assessed on a case by case basis.

4.g Trust in online markets

Introduction

4.245 In general, consumers do not like price discrimination, even if they benefit from it. In online markets, as discussed above, the existence and form of price discrimination is rarely transparent to the consumer. Where there exist information asymmetries such as this, then trust between the consumer and the firm can be an important consideration.

4.246 We first describe some of the empirical evidence on price discrimination and trust. We then briefly discuss how this affects the conclusions discussed above.

⁵⁰ In the models where naïve consumers are better off either (i) there is a mixed population and the firm is able to segment sophisticated and naïve consumers or (ii) the population is entirely naïve or sophisticated, in which case the firm can segment by default.

Empirical evidence on trust and price discrimination

- 4.247 Pavlou, Liang and Xue (2007) undertake an empirical analysis of consumer attitudes to online retailing and note 'the significant negative impact of perceived uncertainty on purchase intentions'. They highlight lack of trust as 'a major impediment to online exchange relationships'.
- 4.248 Garbarino and Lee (2003) experimentally test the effect of personalised pricing on consumer trust. Consistent with previous literature⁵¹ they separate the concept of trust in two dimensions; 'benevolence' trust (holding consumer's interest ahead of self interest) and 'competence' trust (fulfilling the promised service in a reliable manner). Overall trust is an additive function of 'benevolence' and 'competence' trust.
- 4.249 By way of a controlled experiment, they compare survey responses for subjects who were exposed to personalised pricing to responses from a control group; the survey is designed to test both dimensions of trust mentioned above.
- 4.250 The results show that following a personalised pricing experience, subjects express less trust towards online retailers (especially less 'benevolence' trust) and give more weight to 'benevolence' trust in overall trust formation. Interestingly, personalised pricing causes a loss of trust for subjects experiencing higher prices as well as those benefitting from lower prices.

Trust and price discrimination in theory

- 4.251 If consumers lose trust in a market, then they will reduce their demand accordingly. This inward shift in demand causes potential gains to trade to go unexploited, and reduces consumer surplus and profits. We briefly consider the impact of such a reduction in demand in a monopoly and oligopoly setting.
- 4.252 As explained in Section 4.a above, firms price discriminate best-response asymmetrically even though it makes them worse off. They do this because it is individually optimal for a firm to price discriminate given the pricing strategy of its competitors, meaning they end up in a prisoner's dilemma.

⁵¹ Ganesan & Hess (1997), Selnes & Gronhaug (2000), Singh & Sirdeshmukh (2000).

4.253 We note that a significant reduction in demand may mean that it is no longer individually optimal for a firm to price discriminate, meaning that it does not end up in a prisoner's dilemma, and instead an equilibrium in which no firm price discriminates is attained.

Monopoly

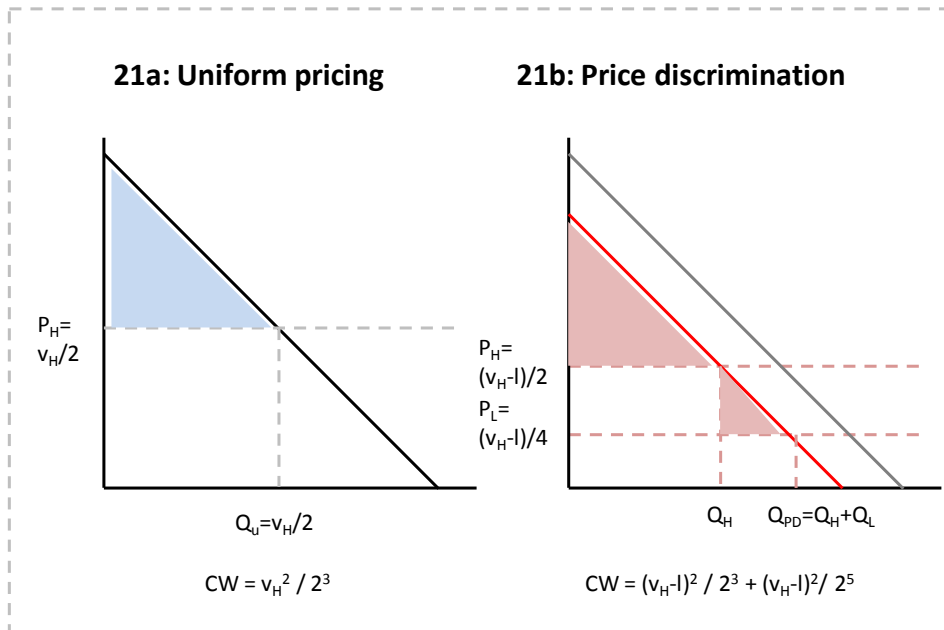
4.254 We now examine the welfare effects of employing personalised pricing using the framework set out in section 4.b.; consumer valuations are uniformly distributed between 0 and VH and marginal costs are normalized to 0.

4.255 In section 4b we found that consumer welfare can increase if the monopolist is able to identify low valuation consumers and tailor a price to their valuation. We wish to examine if this result still holds when we account for a fall in demand triggered by the use of personalised pricing.

4.256 In order to do that, we model the loss of trust as a linear reduction in demand so that consumers' valuation is now distributed between 0 and $(VH-I)$. This is a basic model that assumes that the implied cost to consumers of the loss in trust is constant across all consumers, meaning demand shifts inwards. If the cost is proportional to WTP then the demand would pivot inwards instead.

4.257 We assume that the monopolist is able to segment his customer base into high valuation and low valuation group, after the fall in demand took place, as set out in the figure below. Again, alternative assumptions regarding how firms can price discriminate would be possible.

Figure 21: The effect of a decrease in trust



4.258 Thus the monopolist is able to segment consumers into a high value and low value group, but only at the cost of a reduction in demand.

4.259 The reduction in demand clearly reduces consumer surplus in comparison to price discrimination without a reduction in demand. Price discrimination may still be beneficial to consumers, however, even with the reduction in demand.

4.260 It is possible to derive the critical value of the fall in demand (that is, the critical value of l) at which point discriminatory consumer surplus would be less than uniform consumer surplus. Interestingly, this is the same point at which discriminatory profit becomes less than uniform profit.

4.261 Thus, in this model at least, if the reduction in demand is so large that consumer surplus would be lower than under uniform prices, then in any case the monopolist would itself choose not to price discriminate.

4.262 In general, however, if consumers lose significant trust in online markets then consumer surplus will be reduced, all other things being equal.

Oligopoly

4.263 We describe in Section 4.a above how oligopoly price discrimination is likely to be beneficial to consumers (under best-response symmetry), or at worst the effect is neutral (under best-response symmetry). However, as noted by Garbarino and Lee (2003), in experimental settings consumers do not like price discrimination even if they benefit from it. Thus there may still be a reduction in trust even in oligopoly settings in which price discrimination makes consumers better off.

4.264 The effect of this reduction in demand would, holding all other factors constant, reduce consumer surplus, as potential gains to trade go unexploited. We have not, however, found any models of price discrimination that consider this explicitly.

4.265 If a consumer is price discriminated against online, that may reduce trust in online markets generally, as pricing policies are generally unobservable to the consumers. That is, the consumer reduces trust in any online transaction, as opposed to in that retailer alone. This could impact consumer surplus outside of the market being examined, potentially causing wider harm.

Summary

4.266 A reduction in demand through a loss of consumer trust is likely to reduce consumer surplus in both monopoly and oligopoly settings. This is particularly the case if trust is reduced in online markets generally, instead of just in the firm that is price discriminating.

4.267 It is not straightforward to identify when one might expect the existence of price discrimination to significantly reduce demand. Consumer surveys may usefully give an indication of how consumers feel about particular instances of price discrimination.

5 CONCLUSIONS AND APPLICATION TO ONLINE MARKETS

Introduction

- 5.1 In this section, we summarise the main theoretical results discussed above and attempt to apply some of the conclusions to online markets.
- 5.2 As explained in Section 2, our approach is not to identify a single model that best explains online markets. Instead, we set up a framework indicating circumstances under which we may be concerned about price discrimination.
- 5.3 Much of the literature that we have reviewed does not specifically relate to online retailing, but to general retailing. We have tried to apply this literature to specific features of online markets, but without placing restrictions on the typical form of online retailing. The majority of our framework, therefore, is general and can be applied equally to online or offline retailing.

Framework

- 5.4 By framework, we mean a set of observable features that indicate whether or not price discrimination is likely to be harmful to consumers.
- 5.5 Note that none of these features is intended to be entirely stand-alone. Instead, the observable features provide a guide as to whether price discrimination may, on balance, harm consumer welfare, and should therefore be investigated further.
- 5.6 In broad terms, our proposed framework follows these steps:
- First, identify the intensity of competition in the market.
 - Second, for monopoly and oligopoly in turn, consider the observable features that indicate whether price discrimination is likely to be harmful.

1. Competition

- 5.7 The effect of price discrimination differs significantly depending on whether or not there is competition in the market.
- 5.8 The level of competition in the market will depend on, among other things, (i) the number of firms in the market and (ii) the ease of entry. Each of these should generally be observable in practice.

2. Monopoly

- 5.9 If there is a single firm in the market, then price discrimination can still increase consumer surplus through the Output Expansion Effect. However, we show above that the Appropriation Effect will dominate (meaning price discrimination is likely to be harmful) if the price discrimination (i) is relatively sophisticated and segments consumers into many groups or (ii) identifies high-value consumers rather than low-value.
- **Level of segmentation:** We show above that as the sophistication of the price discrimination increases the Surplus Appropriation Effect becomes more important relative to the Output Expansion Effect. For example, if firms are segmenting consumers into four groups or more, then consumer surplus is likely to be lower than under uniform pricing. Personalised pricing (which we define as sophisticated third-degree price discrimination) by a monopolist is therefore likely to be harmful to consumers.
 - **Type of consumer identified:** we show above that if the firm is simply segmenting its high WTP consumers, as opposed to identifying low WTP consumers, then price discrimination is likely to decrease consumer surplus.
- 5.10 If there are two or more firms in the market, then price discrimination has a different effect, as we discuss below.

3. Oligopoly

- 5.11 The results for oligopoly price discrimination depend on whether the heterogeneity observed is best-response asymmetric (such as where firms

are identifying the brand preferences of consumers) or best-response symmetric (such as where firms are identifying the choosiness or level of switching costs of consumers). In practice, it should be possible to observe what heterogeneity is being conditioned upon and identify whether it is best-response symmetric or asymmetric.

5.12 Under best-response asymmetry, as we discuss above theoretically discriminatory prices could be higher or lower than uniform prices for all consumers. We show above that for most standard demand curves and market sizes best-response asymmetric discrimination is likely to reduce prices. This intuition can be evaluated in practice if it is possible to observe (i) the form of demand faced by the firms or (ii) the relative importance of the strong and weak markets to each firm.

5.13 Under best-response symmetry discriminatory prices will be higher than uniform prices for some consumers and lower for others. However, the aggregate (net) effect is that aggregate weighted prices are no higher, and consumer surplus is unchanged (when each consumer's surplus is equally valued).

5.14 The overall effect of oligopoly pricing is therefore at best beneficial to consumer surplus and at worst neutral (in aggregate). While this is a fairly general result, there are models in which these results do not hold. We illustrate this by considering the following three issues:

- the sophistication of consumers
- the cost of price discrimination, and
- a loss of trust in the market due to price discrimination.

3. a. Sophistication of consumers

5.15 We discuss above how the sophistication of consumers is important, where by sophistication we mean the extent to which consumers are able to rationally optimise over all their choices. This is particularly true in dynamic games where a consumer's purchases today affect the price he is offered tomorrow. Sophistication is also likely to be important in static interactions, although we have not identified models that specifically examine this.

- 5.16 If consumers are sophisticated then by and large the general result that oligopoly price discrimination is likely to be beneficial still holds (although there are exceptions).
- 5.17 If consumers are naïve, then it may be that the firms can exploit that naiveté and make the consumer worse off. This is not always the case, as sometimes the presence of competition and sophisticated consumers can protect naïve consumers.
- 5.18 Where there are different groups of consumers which are either naïve or sophisticated, then the effect on aggregate consumer surplus depends on the relative proportions of each type.
- 5.19 The general indication from experimental evidence is that consumers are not particularly sophisticated, at least not in the way required by the price discrimination literature. It does not follow, however, that they are necessarily entirely naïve.
- 5.20 It is not straightforward in practice to observe whether consumers are naïve or sophisticated in particular cases. It is possible, however, to observe (i) the complexity of the price discrimination and (ii) its transparency. As rough rules of thumb, if the price discrimination is particularly complex or opaque then consumers are less likely to be sophisticated with respect to that form of price discrimination.

3.b. Cost to firms of price discrimination

- 5.21 If firms incur significant costs to price discriminate then in broad terms the prices faced by consumers are likely to be higher. While it does not follow that costly price discrimination is certain to be harmful to consumers, all costly price discrimination is more likely to be so, other things being equal.
- 5.22 It may or may not be possible in particular investigations to observe the monetary and time cost of collecting information for price discrimination.

3.c. Loss of market trust

- 5.23 We describe above how consumers, in general, do not like price discrimination, even if in many cases it may make them better off. This can

manifest itself in a loss of trust in the market, leading to a reduction in demand. This is inefficient and damaging to consumer welfare, as consumers are leaving some beneficial gains to trade unexploited.

5.24 It is possible that trust may be reduced in other online retailers, as well in the retailer that is discriminating, or even that trust is reduced in all online retailing. In this case, the potential harm is much wider.

5.25 While trust is a difficult concept to measure empirically, in practise consumer surveys may give an indication of how strongly consumers feel about price discrimination, and what effect it has on their trust in that particular retailer and in online markets more generally.

Summary

5.26 We summarise below each of the main issues we discuss above, as well as (i) any observable features that could be identified in an investigation and (ii) the likely direction of the effect on how price discrimination impacts on consumer surplus.

Table 4: Summary of framework for evaluating price discrimination

Issue	Observable feature	Effect of price discrimination on consumer surplus
1. Level of competition	- number of firms in market - barriers to entry.	
2. Monopoly	- complexity of discrimination - type of consumer being identified.	Likely to be harmful if sophisticated or identifying high WTP consumers
3. Oligopoly	- best-response symmetry - best-response asymmetry.	Neutral under best-response symmetry, beneficial under asymmetry, subject to other issues below
3.a. Sophistication of consumer	- complexity of price discrimination - transparency.	More likely to be harmful if complex or not transparent
3.b. Cost of price discrimination	- source of information being collected - cost to the firm of	More likely to be harmful if costly

	collection.	
3.c. Loss of market trust	- reduction in demand caused by introduction of price discrimination	More likely to be harmful if reduces trust in the whole online market and demand

Characteristics outside the framework

5.27 The framework above summarises characteristics for which we have been able to identify a general (albeit, in some cases, approximate) relationship. We have also identified various key characteristics for which we have been unable to identify a general relationship. These may still be relevant in particular investigations, depending on the information being observed:

- **Accuracy of observed heterogeneity:**⁵² in practice, an online monopolist is likely to make errors when estimating whether a consumer has high or low WTP, either by mischaracterising (i) low WTP consumers as high WTP or (ii) high WTP as low WTP. All other things being equal, this form of inaccuracy will be beneficial to consumers if the latter effect dominates the former, and vice versa. In practice, however, the true WTP is unlikely to be observable in any investigation, which explains why this characteristic is not part of the framework.
- **Advertising:**⁵³ There is evidence that firms are targeting their advertising to specific consumers, based on heterogeneity they observe online. This form of advertising can be beneficial to consumers, as they are more likely to see relevant adverts than irrelevant adverts, which helps the matching process between firms and consumers. The provision of such advertising may also provide revenue support to the online services the consumer is using such that the consumer is not charged directly for these. However, it may also be harmful when combined with price discrimination, as it can increase the perceived level of differentiation between firms (effectively increasing their market power). It is not possible, therefore, to draw a general relationship between the effect of price discrimination and advertising.
- **Avoidance by consumers:**⁵⁴ The ability to avoid price discrimination by, for example, anonymising their online profile increases the strategy

⁵² See paragraph 0.

⁵³ See Section 4.d above.

⁵⁴ See paragraph 0 above.

space of consumers, which may increase their consumer surplus (in simple terms, they have more options available to them and so may be able to obtain a better outcome). However, this is not always the case, for two reasons. First, consumers may still have an incentive to anonymise themselves if price discrimination is generally beneficial to consumers (such as under best-response asymmetry), as consumers are generally better off if they are (or can pretend to be) indifferent between the competing firms. If all consumers did this, however, price discrimination would break down and the firm would return to uniform pricing, to the detriment of all consumers (in other words, consumers end up in a Prisoner's Dilemma). Second, anonymisation may in fact be observable to firms (consumers have a blank online profile, for example), and may in fact be an indicator of high WTP (as those with low WTP have less to lose from price discrimination and would not incur the cost of anonymisation). It is not possible, therefore, to draw a general relationship between the effect of price discrimination and anonymisation.

Conclusion

- 5.28 Online personalised pricing is likely to take various different forms regarding, among others, (i) the level of competition in the market, (ii) the complexity of price discrimination (and thus how likely consumers are to understand it) and (iii) the cost to the firm of discriminating. It is not possible, therefore, to conclude whether, in general, online personalised pricing is harmful or beneficial to consumers.
- 5.29 Instead, in specific investigations it will be necessary to follow the framework we summarise above, based on the specific factors of the case.
- 5.30 We note, however, that many of the market characteristics that make price discrimination more likely to be harmful apply to online retailing more than they do to offline retailing. In particular, online retailing is (i) more likely to be complex, as firms can observe more heterogeneity and (ii) more likely to be opaque, as consumers may not be able to observe the price being quoted to other consumers. This does not mean that personalised pricing online, as a rule, is likely to be harmful: it just means that caution should be taken when applying this framework to individual cases of price discrimination.

6 BIBLIOGRAPHY

- Acquisti, Alessandro, and Hal R. Varian (2005). 'Conditioning prices on purchase history.' *Marketing Science* 24.3: 367-381.
- Armstrong, M. (2006): 'Recent Developments in the Economics of Price Discrimination,' in Blundell, R. and Newey, W.K. and Persson, T., (eds.) **Advances in Economics and Econometrics: Theory and Applications: Ninth World Congress: volume II.** (pp. 97-141). Cambridge University Press: Cambridge, UK.
- Armstrong, M., and J. Vickers (2001): 'Competitive Price Discrimination,' *Rand Journal of Economics*, 32(4), 579–605.
- Bester, H., and E. Petrakis (1996): 'Coupons and Oligopolistic Price Discrimination,' *International Journal of Industrial Organization*, 14(2), 227–242.
- Branca Esteves, R. (2009): 'A Survey on the Economics of Behaviour-Based Price Discrimination', NIPE working paper.
- Caminal, R., and C. Matutes (1990): 'Endogenous Switching Costs in a Duopoly Model', *International Journal of Industrial Organization*, 8(3), 353–373.
- Charness, Gary, and Dan Levin (2003). 'When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect.' *Complexity, and Affect* (November 17, 2004).
- Chen (2006): 'Marketing Innovation,' *Journal of Economics and Management Strategy*, 15(1), 101–123.
- Chen , Y. (1997): 'Paying Customers to Switch,' *Journal of Economics and Management Strategy*, 6(4), 877–897.
- Chen, Y. (2005), *Oligopoly Price Discrimination by Purchase History, The Pros and Cons of Price Discrimination*, The Swedish Competition Authority, Stockholm, 101-130.
- Chen, Y., Narasimhan, C. and Zhang, Z. (2001), *Individual Marketing with Imperfect Targetability*, *Marketing Science* 20, 23-41.
- Coase, Ronald H. (1972), 'Durability and monopoly.' *JL & Econ.* 15: 143.
- Conitzer, Vincent, Curtis R. Taylor, and Liad Wagman (2012). 'Hide and seek: Costly consumer privacy in a market with repeat purchases.' *Marketing Science* 31.2: 277-292.

- Corts, K. (1998): 'Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment,' *Rand Journal of Economics*, 29(2), 306–323.
- De Corniere, Alexandre and Nijs, Romain De, (2012) *Online Advertising and Privacy*. Available at SSRN.
- Esteves, Rosa-Branca (2012), 'Price Discrimination with Private and Imperfect Information', NIPE WP 12/ 20.
- Fudenberg, D., and J. Tirole (2000): 'Customer Poaching and Brand Switching,' *Rand Journal of Economics*, 31(4), 634–657.
- Ganesan, Shankar, and Ron Hess (1997). 'Dimensions and levels of trust: implications for commitment to a relationship.' *Marketing letters* 8.4: 439-448.
- Garbarino, Ellen, and Olivia F. Lee (2003). 'Dynamic pricing in internet retail: effects on consumer trust.' *Psychology & Marketing* 20.6: 495-513.
- Gehring, T., Shy, O and Stenbacka, R. (2011), *History-Based Price Discrimination and Entry in Markets with Switching Costs*, *European Economic Review* 55, 732-739.
- Hoffmann, Florian, Roman Inderst, and Marco Ottaviani (2012). 'Tailored Advertising and Consumer Privacy.'
- Holmes, Thomas J. (1989) 'The effects of third-degree price discrimination in oligopoly.' *The American Economic Review* 79.1: 244-250.
- Iyer, G., Soberman, D., & Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3), 461-476.
- Liu, Q. and Serfes, K. (2004), *Quality of Information and Oligopolistic Price Discrimination*, *Journal of Economics and Management Strategy* 13, 671-702.
- OFT (2009) 'Advertising of Prices'.
- OFT (2010) 'Online Targeting of Advertising and Prices'.
- OFT (2013), 'Personalised Pricing'.
- Pavlou, Paul A., Huigang Liang, and Yajiong Xue (2007). 'Understanding and mitigating uncertainty in online exchange relationships: A principal-agent perspective.' *Mis Quarterly* 31.1: 105-136.
- Pigou, A. (1932). *The Economics of Welfare*, 1920. McMillan&Co., London.

- Robinson, J. (1933) *The Economics of Imperfect Competition*, London: Macmillan.
- Rochet, J.-C., and L. Stole (2002): 'Nonlinear Pricing with Random Participation,' *Review of Economic Studies*, 69(1), 277–311.
- Salop, S., & Stiglitz, J. (1977). Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies*, 44(3), 493-510.
- Selnes, F., & Grønhaug, K. (2000). Effects of supplier reliability & benevolence in business marketing. *Journal of Business Research*, 49, 259–271.
- Shin, J., & Sudhir, K. (2010). A customer management dilemma: When is it profitable to reward one's own customers?. *Marketing Science*, 29(4), 671-689.
- Singh, J., & Sirdeshmukh, D. (2000). Agency and trust mechanisms in consumer satisfaction and loyalty judgments. *Journal of the Academy of Marketing Science*, 28(1), 150-167.
- Stigler, G. (1987), *A Theory of Price*, New York: MacMillan.
- Stole, L. (2006): 'Price Discrimination and Imperfect Competition,' in *Handbook of Industrial Organization: Volume III*, ed. by M. Armstrong, and R. Porter. North-Holland, Amsterdam.
- Taylor, C. R. (2003). Supplier Surfing: Price-Discrimination in Markets with Repeat Purchases. *RAND Journal of Economics*, 34(1), 223-46.
- Thisse, J.-F., and X. Vives (1988): 'On the Strategic Choice of Spatial Price Policy,' *American Economic Review*, 78(1), 122–137.
- Varian, H. (1985): 'Price Discrimination and Social Welfare,' *American Economic Review*, 75: 870-875.
- Weizsäcker, G. (2010). Do we follow others when we should? A simple test of rational expectations. *The American Economic Review*, 2340-2360.