Essays in Industrial Organization

Kristofer Hammarbäck

Regent's Park College

A thesis submitted for the degree of
Doctor of Philosophy in Economics

Hilary Term 2018

Word count: Approximately 44,000 words, taking p. 63 as typical page.

Statement of authorship: All parts of this thesis are solo-authored and have not been submitted as part of fulfilment of requirements for a degree in this university or elsewhere, with the exception of Chapter 2 which is based on my M.Phil. thesis.

Contents

Abstract **4**

Acknowledgements **5**

1 Introduction **6**

2 Unawareness and Communication in Vertical Relationships **10**
- 2.1 Introduction 10
- 2.2 Model 15
- 2.3 Observable Effort 20
- 2.4 Unobservable Effort 29
- 2.5 Welfare Analysis 41
- 2.6 Conclusion 44

3 Markets with Differentiated Add-on Products **46**
- 3.1 Introduction 46
- 3.2 Model 51
- 3.3 Random Selection 53
- 3.4 Endogenous Selection 65
- 3.5 Conclusion 75
- 3.A Appendix 77

4 Competition and Bundling: An Empirical Analysis of European Telecom Markets **88**
- 4.1 Introduction 88
- 4.2 European Telecom Markets 94
- 4.3 Data Description 95
- 4.4 Model 101
- 4.5 Estimation 104
- 4.6 Results 112
- 4.7 Counterfactual Simulations 116
- 4.8 Conclusion 122
- 4.A Appendix 124

Bibliography **130**

Abstract

This thesis contains three main chapters, each providing analysis into topics in the area of industrial organization.

Unawareness and Communication in Vertical Relationships

I analyse communication in a setting where a distributor offers a contract to a producer over the supply of a product. Both firms can exert effort, which may be observable or unobservable. The distributor is aware of a pay-off relevant contingency and can choose to communicate this to the unaware producer. The distributor may in communicating face a trade-off between inducing optimal effort inputs and lowering the cost of ensuring participation. Communication occurs more often when effort is unobservable as the distributor is less able to benefit from his information. Joint surplus of the firms is higher when effort is unobservable than when effort is observable.

Markets with Differentiated Add-on Products

I analyse a duopoly market where firms supply differentiated core and add-on products. Some consumers are uninformed about preferences and prices of add-on products. I find that firms will price add-ons above marginal cost even with fully informed consumers and that add-on profits are not competed away in core market competition. The market outcome leads to inefficient matching between firms and consumers who purchase too few add-ons. The presence of uninformed consumers is harmful for informed consumers because of higher add-on prices. If consumers can decide whether to be informed or not, add-on prices rise with the fraction of informed consumers.

Competition and Bundling: An Empirical Analysis of European Telecom Markets

Using data on European telecom markets, I estimate a structural model of demand for mobile and broadband services. I use two data sources that provide choice information at different aggregation levels which the model is constructed to account for. The estimated model is used to conduct merger simulations. Issues focused on are the importance of pricing complementarities arising between firms' standalone product offers as well as the level of consumer demand for bundling. The simulations imply that merger outcomes can significantly depend on complementarity effects. I also find that higher demand for bundling leads to higher price increases from mergers, as firms become stronger substitutes.

Acknowledgements

I would like to thank my supervisors Mark Armstrong, Howard Smith and John Quah for providing invaluable insight, advice and support throughout the writing of this thesis.

I am grateful for comments on my work from my confirmation of status assessor Simon Cowan, and for seminar comments from Danny McGowan and Meg Meyer. I would like to thank my colleagues Gregorio Curello, Pawel Dziewulski, Claudia Herresthal, Sam Jindani, Josh Lanier, Anna Sanktjohanser and Segye Shin for comments and discussions that have helped shape the chapters of this thesis. I am also grateful for numerous informal discussions that have helped develop my ideas; in particular I thank Maximilian Igl and Patrick McGrath for suffering through many lengthy meal conversations about my work. For market insights regarding my final chapter, I thank Rupert Wood.

I gratefully acknowledge the Wallenberg foundations and the Department of Economics at the University of Oxford for financial support, as well as the University of Oxford Advanced Research Computing (ARC) for use of their facilities.

Finally, I would like to extend my gratitude to all family and friends in Oxford and elsewhere who have provided me with invaluable support throughout my doctoral studies.

1 Introduction

This thesis contributes to the literature by providing three novel pieces of analysis that all relate to industrial organization. Therefore, the thesis is organised into three main chapters. In the first main chapter, I consider communication incentives in a supply chain where one of the parties to a contract is unaware of a payoff-relevant contingency. The subsequent chapter analyses competition in add-on product markets where the add-ons are differentiated, with a focus on the role of consumer information. The final chapter analyses the implications of recent bundling trends in telecom markets for competition policy by empirically estimating a structural model of telecom demand. The three chapters of the thesis are all self-contained and standalone papers. As such, the chapters are here introduced separately in turn.

In chapter 2, I analyse firm behaviour in the context of vertical contracting with informational asymmetries. The chapter posits a model where a manufacturer and a distributor are contracting about the supply of a product to a downstream market where both parties can exert effort to improve the revenue raised by the product. In the situation considered, the contracting parties have an asymmetry of awareness regarding a potentially payoff-relevant contingency. That is, there is some event affecting payoffs that may occur, but only the distributor is aware of the possibility of this event occurring. In this setting, the chapter derives predictions about contractual features under asymmetric awareness. Using these predictions, I answer the question of under which circumstances the distributor has incentives to communicate information to expand the producer's awareness. The analysis makes a distinction between producer effort input being observable or unobservable to the distributor and compares outcomes arising from the two cases.

I find that in choosing whether to reveal information or not, the distributor faces a trade-off between inducing optimal effort provision and a potential ability to pay the producer less than the reservation payment. Furthermore, the imposition of a monotonicity requirement on contract forms creates an asymmetry between information that is 'good' or 'bad'. Due to this requirement, the distributor has incentives to offer a high-powered contract when there is good news to be told and offer a flat repayment contract when there is bad news to be told. This means that the model provides strict predictions about what contracts the distributor will offer, even if efforts are observable and firms are risk neutral, which in a standard setting would imply a multiplicity of optimal contracts. Because of these incentives, it is found that if there is good news to be told, the distributor always wishes to inform the producer because the distributor becomes unable to exploit his superior awareness. With bad news, however, the distributor can design a contract that pays less than the producer is believed to be paid. Therefore, it might become optimal for the distributor to conceal the event, despite this leading to suboptimal effort inputs.

The obtained results lead to a striking welfare result; that it is ex ante better for the supply chain if the distributor is not able to observe the producer's effort input. The source of this result lies in that communication is more likely to occur when effort is unobservable and therefore more often leads to the jointly optimal effort inputs. There is no trade-off with agency costs since players are risk neutral and so unobservable effort is welfare enhancing for the supply chain. There is also a tension in the firms' preferences over effort observability, in the sense that the distributor is better off when effort is observable while the producer is better off when effort is unobservable.

In chapter 3, I analyse the effectiveness of competition in markets with add-on products. In particular, I focus on the case where add-ons are differentiated. For instance, the content of the add-on may be different at each firm and consumers may therefore have varying preferences regarding which add-on product they prefer. Furthermore, as many add-on markets are characterised by uninformed consumers, I incorporate this into the analysis by letting some consumers be uninformed about add-on content and prices. My focus is on analysing the equilibrium properties that arise when add-ons are differentiated and whether the market outcome under competition is socially efficient. As such, I construct a duopoly model where firms are differentiated on a Hotelling line, but where add-on products are also differentiated in the sense that consumers may demand the add-on product at one firm but not at the other. In a second version of the model, I analyse the case where consumers rationally decide whether to become informed about the add-ons or not.

I find that firms set add-on prices above marginal costs even with a fully informed consumer population, as a result of the add-on differentiation. Firms also make more profits when add-ons exist than when they do not, in contrast to a profit irrelevance result of add-ons that applies when add-ons are homogenous. Firms do compete away parts of the add-on profits in the core product competition, but incompletely so. These results stem from that with differentiated add-ons, a firm's inframarginal consumers are on average more profitable than a firm's marginal consumers. Therefore, incentives to compete for marginal consumers are lower than with homogenous add-ons. In the case where consumers rationally decide whether to become informed, I find the striking result that add-on prices rise when more consumers are informed of add-ons. This result rests on a price discrimination effect that arises due to the selection of consumers that are informed being those with weakest core product preferences.

The welfare analysis indicates that the market equilibrium is inefficient in the sense that consumers are not efficiently matched to firms and end up purchasing too few add-ons. In the basic model, welfare therefore improves when more consumers become informed. This is both because newly informed consumers make better choices and because more informed consumers drive down add-on prices, which improves efficiency. The analysis of the model where consumers decide whether to become informed or not reveals that there will be too few informed consumers in equilibrium from a social welfare perspective since consumers do not internalise the positive effect their information have on firm profits. However, from a consumer welfare perspective, there is overinvestment in information acquisition, which derives from the fact that prices are then increasing in consumer information. The results imply that information disclosure interventions are beneficial, as better information improves matching between firms and consumers, but also highlight that it may be important to consider what selection of consumers it is that will become informed as this can have a substantial impact on the resulting outcomes.

In chapter 4, I undertake an empirical analysis of competition and bundling in European telecom markets. There has recently been a strong trend in telecom markets towards bundling of products that were previously only sold separately. This trend has been accompanied by a wave of consolidation both within and across mobile and broadband providers, raising interest in competition overview of telecom markets. I investigate the implications of the increased prominence of bundling of mobile and fixed services for competition policy by analysing counterfactual mergers under various environments. To conduct this analysis, I use data at both the firm and the household level to estimate a structural model of demand for mobile and broadband services. The model allows households to have a preference for bundling and firms to set bundling discounts. The firm level data provides information on a firm-product specific level, while the household level data only provides choices at the product group level, so the model is adjusted to allow for the data sets providing choice information at different aggregation levels. The price coefficient is estimated using a supply-side assumption on pricing behaviour as well as additional supply-side information.

The findings imply that the level of pricing complementarities can be important for the price effects of mergers. These complementarities arise as some consumers that are initially purchasing two standalone offerings may switch to a bundle following a price increase in one of these standalone products, therefore reducing demand for the other standalone product. By comparing simulation outcomes of two illustrative mergers, it is found that the resulting price increases are 6.31 percentage points higher in the merger where the firms are less complementary, which translates to a difference in impact on consumers equivalent to $4.90 of surplus per consumer per month. Furthermore, markets with higher levels of bundling exhibit stronger reduction of competition as a result of mergers, as firms are more strictly substitutable when bundles are prominent. This intuition is illustrated by merger simulations where mergers are conducted for actual market conditions as well as counterfactual conditions, where demand for bundles is higher, and resulting outcomes compared. The results indicate that if market shares for bundles in markets with lower bundle demand would be at the level of the market with the highest share of bundles, price increases resulting from mergers are 1.79 percentage points higher. This is equivalent to a difference in impact on consumers of $1.61 of surplus per consumer per month. Implications for policy are that simple market share analysis may be insufficient to conclude pricing effects from a merger with mixed bundling and that one may expect mergers to be more anti-competitive the more prominent bundles are in the market. These results have implications for the strictness of merger control in the light of recent telecom market trends, and for other markets that are significantly characterised by bundling.

2 Unawareness and Communication in Vertical Relationships

2.1 Introduction

When a distributor writes a contract with a producer there is a possibility that one party to the contract is more informed than the other party regarding some relevant piece of information. For instance, the distributor may possess better information about local market conditions and be more informed regarding the key determinants of local demand. This type of situation has been frequently analysed in economics using standard models of asymmetric information. In such models, the uninformed party does not know what the state of the world is but knows what the possible states are and forms a belief over the probability of each state occurring. One kind of situation which standard models leave out, however, is where the less informed party to the contract may not know that there is something to be known. That is, the uninformed party may not be *aware* that the other party is superiorly informed. Unawareness relates to the possibility that an agent is unable to take some fact(s) into account in his decision-making, which could be a result of cognitive or informational limitations.

This chapter considers the concept of unawareness as applied in a setting where a distributor is contracting with a producer. In particular, the chapter investigates whether it is in the interest of a superiorly aware distributor to expand the awareness of an unaware producer. In doing so, implications for contracting under asymmetric awareness is derived. Supply chain contracting is an appropriate area to study asymmetry of awareness, as the parties may have differing levels of awareness arising from their differing vertical positions in the chain. For example, in the supply chain of a consumer product, the distributor of the product may be highly aware of the contingencies that determine how well the product will perform on the consumer market. The producer of the product, specialising in production rather than distribution, may be less aware of the underlying determinants of downstream market success. Because of bounded rationality, information acquisition costs or simply the given information structure the producer may not just be uninformed of the relevant determinants, but unaware of some relevant determinants and fail to take them into account in decision making.

An interesting question to ask is whether unawareness in this context would be robust in the sense that if the distributor was superiorly aware, would the distributor have incentives to expand the awareness of the producer by sharing the information? As such, this chapter investigates in a contracting model whether the superiorly aware distributor does have such incentives. A second aim of the chapter is to derive insights into predictions of contracting outcomes under asymmetry of awareness. These issues are analysed in a model of contracting in a supply chain where the more aware party is given the opportunity to inform the less aware party prior to a contract being signed. Contracts take the form of a monetary transfer that is contingent on realised revenues. Both parties are able to exert effort into improving the profitability of the good on the final market. These efforts may be interpreted as advertising efforts or efforts toward improving the quality or reliability of the product. Importantly, the chapter will make a distinction between efforts being unobservable and observable. In the quality interpretation, this may refer to the extent to which the quality of the product is observable before being brought to the market.

The definition of unawareness employed in this chapter is based on Modica and Rustichini (1999). When the producer is unaware of a contingency, this will mean that she is unaware that she is unaware of that contingency and so on *ad infinitum*. The behavioural implication is that whenever the producer is unaware, she will act as if the environment that she perceives is the correct description of the world. The producer is not able to reason about the possibility that the distributor may be aware of some contingencies that she herself is not aware of and does not realise that absence of communication is, in fact, an action deliberately chosen by the distributor. In contrast, the distributor fully comprehends all aspects of the situation, including the producer's limited awareness. This asymmetry of awareness could as mentioned be a result of bounded rationality in the decision making or information processing on the part of the producer, preventing the producer from reasoning about aspects of the environment that she may be unaware of.

The results of the model provide insights into contractual design under asymmetric awareness and mutual risk neutrality. For example, in standard models with observable effort, there are multiple contract shapes that are optimal, as only expected payment matters and awareness is symmetric. Under asymmetric awareness, the analysis in this chapter demonstrates that there is a unique contract shape that is optimal. This is because even if multiple contracts seem optimal from the producer's limited awareness, there is a unique contract that minimises the expected transfer the distributor makes. Also, in standard models with mutual risk neutrality, the principal who offers the contract is indifferent to whether effort is observable or not as the project could be 'sold' to the agent in the unobservable effort case. However, in the setting with asymmetric awareness, there are instances in which the distributor strictly prefers effort to be observable. This is because the uniquely optimal contract under observability may no longer be feasible in the unobservable effort case as it would cause suspicion on the part of the producer. A key assumption driving these results is that I assume contracts to be non-decreasing, motivated by the difficulty of justifying transfer schedules that are decreasing in revenues raised.

Regarding communication incentives, it turns out to be convenient to make a distinction between 'good' and 'bad' information to be shared. Under the definitions used and the framework of the model, I find that good news will always be communicated while bad news may not be. This asymmetry is sourced from the fact that when there is good news to be told, the distributor is unable to exploit the fact that he has superior awareness. With hidden good news, he cannot design a contract that in expectation pays less than the producer believes to be paid and would still be accepted. Revealing the information leads to optimal effort choices, so it must then always be optimal for the distributor to inform the producer. When there is bad news, however, the distributor is able to design a contract where he pays less in expectation than the producer believes to be paid and participation is therefore cheaper to satisfy. This creates a trade-off between paying less to satisfy participation and inducing optimal effort levels. Due to this trade-off, it is sometimes optimal to reveal the unknown contingency and sometimes not to reveal. Finally, I show that there exist instances with bad news where the news is not revealed when effort is observable but is revealed when effort is unobservable, as the distributor's ability to exploit his superior awareness is limited when effort is unobservable. In this sense, there is more communication in the supply chain when effort is unobservable.

The welfare analysis demonstrates that the expected joint surplus for the firms is higher when effort is unobservable compared to the case where effort is observable. This result derives from the insight presented above that unobservable effort means that there is a wider range of parameter values for which information is communicated. Since communication of information always leads to both firms exerting the jointly optimal effort inputs, unobservable effort therefore leads to higher joint surplus for the firms. There is however a tension in the firms' preferences over effort observability. While joint surplus is higher under unobservable effort, the distributor prefers effort to be observable as this makes it more likely that he can more cheaply satisfy the producer's participation. The producer prefers effort to be observable as this guarantees an expected utility equal to the reservation level.

The remaining part of this section describes the relevant literature. Section 2.2 outlines the general model set-up and key assumptions made. Section 2.3 analyses the case of producer effort being observable, while section 2.4 analyses unobservable producer effort. Section 2.5 provides a welfare analysis and section 2.6 concludes.

Literature

A recent contribution related to the present chapter is Auster (2013), who applies unawareness to the canonical principal-agent model. In hers as in my model, the more aware player has the opportunity to expand the unaware player's awareness when making a contract offer. It is then found that the principal making the contract offer faces a trade-off between ease of inducing effort and cost of ensuring participation. Unlike my model, the trade-off in her model depends on the realisation of contingencies being informative of effort and on the agent being risk-averse. In my model, contingencies are independent of effort and agents are risk neutral, so the source of the trade-off is different. Filiz-Ozbay (2012) also analyses a related situation where a superiorly aware insurer offers a contract to a less aware insuree. A main difference in her model is that hard information is not assumed; the less aware agent is allowed to have a wide range of beliefs over the probability of contingencies occurring and these endogenous beliefs become part of the solution concept. This assumption drives the results that focus on the extent of incompleteness of contracts where it is shown that under some refinements there are no equilibria where contracts are complete. Another related model is Tirole (2009) who presents a model where a buyer and a seller are contracting about the delivery of a good/service. Both players are unaware of what may go wrong in the design of the good, but both parties are able to expand their awareness by paying cognitive costs. So in contrast to Auster (2013) and Filiz-Ozbay (2012), in Tirole's model the agents are aware of their unawareness and may act as to reduce this unawareness.

Awareness and unawareness is a topic discussed ever since Hintikka (1962), in his seminal work on epistemic logic, pointed out the logical omniscience problem: in standard knowledge-related work agents know all tautologies and know all logical consequences of their knowledge, which may seem inappropriate for resource-bound agents. Modica and Rustichini (1999) define unawareness by discussing three possible states of an agent's knowledge regarding a proposition p. One state is when the agent knows whether p is true or false while another state is when the agent knows that he does not know whether p is true or not. Standard theory restricts itself to considering these two possible states of knowledge, but this leaves out a third possibility; that an agent does not know p, does not know that he does not know p and so on *ad infinitum.* Modelling of unawareness in economics has in the last decades gained increasing attention, in particular after Dekel *et al.* (1998) showed that standard state space models are unable to incorporate non-trivial unawareness. This result has been followed by developments in generalised state space models that are able to model unawareness, see Li (2009), Heifetz *et al.* (2006) and Heifetz *et al.* (2013). The developments in state spaces capable of incorporating unawareness have sparked various applications in economic theory; for an application in decision theory see Schipper (2013) and for applications in game theory, see Feinberg (2012) and Halpern and Rego (2014).

There is also a literature on asymmetric information in the principal-agent model where the principal has private information and may signal information using the contract offer, as analysed in a general sense by Maskin and Tirole (1990) and Maskin and Tirole (1992). Aghion and Hermalin (1990) analyse contractual signalling in the context of financial contracting, and focus on separating and pooling equilibrium. While related in the sense of asymmetric information with an informed principal, my unawareness assumption removes any signalling aspect of the contract offer. The closest related paper from this literature that has both an informed principal and relationship-specific investments is Vasconcelos (2014).

There is an existing literature on supply-chain contracting with asymmetric information regarding demand conditions. A key difference between my work and this literature is that this literature tends to focus on the efficiency of the quantity supplied through the chain rather than effort incentives for the parties in the chain. Arya and Mittendorf (2004) study how a firm may elicit a distributor's private demand information through a buyback possibility, where a distributor under adverse demand conditions would prefer the buyback contract but a distributor under advantageous demand conditions prefers the non-buyback contract. Liu and Ozer (2010) study information sharing in the supply chain, but instead consider that the supplier may be superiorly informed, and study information sharing in this setting. The contracts are exogenously given, and the supplier is then able to inform the distributor of more precise demand conditions. As such, the model can be interpreted as the players both being unaware at the time of signing the contract that the supplier will become more informed. As in my model, they find that the information sharing decision may not be jointly optimal, which derives from the result that information in their model is shared whenever doing so increases the order quantity.

2.2 Model

The modelled situation involves two players, denoted as the *Producer* (P) and the *Distributor* (D). They are involved in the supply chain of a good; the producer producing the good and the distributor distributing it to the downstream market. Both players can exert effort into improving the quality/desirability of the good, which they do at a cost $c(e_i)$, $i \in \{P, D\}$. The common cost function c has standard properties: $c(0) = 0,\ c'(.) > 0, c''(.) > 0$. The vector of efforts (e_P, e_D) is denoted as e and efforts are restricted to be non-negative, so that $e \geq 0$. The efforts supplied by the two players generate a state-dependent revenue function $\pi(e, s)$ that determines the revenue raised on the downstream market. The state together with efforts fully captures all determinants of the revenue function, such that the revenue function is non-stochastic when conditioned on a state and effort vector. The function is continuous and twice differentiable in both players' efforts.

The primitive objects of the uncertainty of the model are stochastic contingencies. Let the full set of contingencies be Θ with typical object $\theta_i \in \Theta$. While it is possible to extend this framework to numerous contingencies, the formulation used will have only two contingencies, i.e. $\Theta = \{\theta_1, \theta_2\}$. Each θ_1, θ_2 is then a random variable which can take two possible values, 0 or 1. The natural interpretation of this representation is that the contingency either occurs (1) or does not occur (0). Denote $p_i \in (0, 1)$ as the probability of contingency i occurring. The contingencies are assumed to be independent and the realised value of θ_i is denoted by $\tilde{\theta}_i$. A state s is then defined as a collection of realisations of contingencies, $s = \{\tilde{\theta}_1, \tilde{\theta}_2\}$, where the first entry will refer to the realisation of the first contingency and the second entry refer to the realisation of the second contingency. The set of all such possible collections is the state space S, i.e. $S = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$.

To capture the idea of asymmetric awareness, a set-up similar to that used in Auster (2013) is employed, which in turn is based on the work by Heifetz *et al.* (2006). Firstly, I will use the 'normalisation' that if a player is not aware of a contingency, the player will act as if that contingency will not occur. That is, the player will act as if the contingency is necessarily equal to 0. Throughout the analysis, it will be assumed that before any interaction occurs, the distributor will be aware of both contingencies, while the producer will only be aware of θ_1. The distributor's perceived state space is then S as previously defined, while the producer has a limited perceived state space $S^P = \{(0, 0), (1, 0)\}$. The producer is also unaware that the distributor has superior awareness and is therefore initially unaware of the existence of θ_2, while the distributor fully understands the structure of both players' knowledge and awareness.

The order of events in the model is as follows. Before any interaction, the distributor is given the opportunity to inform the producer about the stochastic process determining revenues. This is done by the distributor sending a message m consisting of a subset of contingencies. If a contingency is included in the message, the producer becomes aware of this contingency. The distributor has perfect knowledge about the awareness of the producer, so any contingency that is already known by the producer is included in the message by default. Therefore, the possible messages are $m = \{\theta_1\}$ or $m = \{\theta_1, \theta_2\}$, where the former is denoted as the *no revelation* case and the latter as the *revelation* case. The contingencies contained in a message are verifiable, in the sense that once communicated, both players understand the revenue impact of the occurrence of the contingency, what the probability of the contingency occurring is and this is common knowledge. The message is therefore 'hard information', i.e. the distributor is not able to lie about the effect of a contingency, he can only choose to reveal it or not reveal it.

Once the communication stage is completed, the distributor makes a contract offer to the producer. The producer may either accept or reject the contract and in the case of rejection will receive the outside option $\bar{U}$. If effort is observable, the contract offer consists of a pair $(e, t(\pi))$ which specifies the effort levels for both players and a monetary transfer that is made from the distributor to the producer contingent on realised revenues. If effort is unobservable, the contract is simply a transfer schedule. After a contract has been accepted, efforts for both players are chosen. Finally, the revenue in the downstream market is realised and captured by the distributor, who in turn makes a payment $t(\pi)$ to the producer in accordance with the signed contract. The ordering of events is therefore as follows:

1. The distributor chooses whether to reveal the contingency or not by choosing a set of contingencies that he communicates to the producer.

2. The distributor offers a contract $(e, t(\pi))$ if efforts are observable or a contract $t(\pi)$ if efforts are unobservable.

3. The producer accepts or rejects the contract.

4. Both players choose their effort levels.

5. Revenues are realised and accrued to the distributor, a payment $t(\pi)$ is paid by the distributor to the producer.

Note that the transfer is contingent on the realised revenue and not on the state. Once efforts have been fixed, the players believe that only a discrete number of revenue realisations and therefore also transfer realisations are possible. In fact, fixing efforts at some arbitrary level e', a player under full awareness will believe that there may be four distinct transfer realisations, one for each state in S. Instead, under limited awareness a player believes that only the transfers given by $t(\pi(e', 00))$ and $t(\pi(e', 10))$ are possible.[1] Despite this, the transfer schedule must still be specified for any possible revenue realisation. This imposition is particularly important when effort is unobservable, as by altering effort inputs any revenue realisation may become possible and so the shape of the transfer schedule outside of levels perceived possible

1 The comma in the state will at times be dropped for brevity.

under the contract may be important to determine deviation incentives. Next, there are a number of assumptions made regarding the properties of the revenue function $\pi(e, s)$.

Assumption 1.

(i) $\frac{\partial \pi(e,s)}{\partial e_i} > 0 \ \forall \ i \in \{P, D\}, \ \forall \ s \in S$

(ii) $\frac{\partial^2 \pi(e,s)}{\partial e_i^2} \leq 0 \ \forall \ i \in \{P, D\}, \ \forall \ s \in S$

(iii) $\frac{\partial \pi(e_i, e'_j, s)}{\partial e_i} = \frac{\partial \pi(e_i, e''_j, s)}{\partial e_i}, \ \forall \ i \in \{P, D\}, \ \forall \ e'_j, e''_j \in R_+, \ \forall \ s \in S$

(iv) $\pi(e, s') \neq \pi(e, s''), \ \forall \ s', s'' \in S, \ \forall \ e \in R^2_+$

(v) $\pi(e, 1\tilde{\theta}_2) > \pi(e, 0\tilde{\theta}_2) \ \forall \ \tilde{\theta}_2 \in \{0, 1\}, \ \forall \ e \in R^2_+$

The first two assumptions imply that efforts always have a positive effect on revenues and that the returns are weakly diminishing. The third part, which is equivalent to the revenue function being additively separable in player efforts, implies that players do not need to make conjectures about the other player's effort choice and greatly simplifies the analysis. The fourth assumption implies that no two states imply the same revenue level, which avoids degenerate cases. The last condition implies that the occurrence of the first contingency always positively impacts revenues.

Next, there are a number of assumptions made on the set of feasible transfer schedules. Firstly, consider that there is some revenue level that the producer is unaware of, but the distributor knows to be possible under full awareness. Suppose that this revenue level is higher than any revenue level the producer believes to be possible. Then the distributor would like to propose a contract that sharply drops for higher revenue levels, which the producer is happy to accept since she does not believe that such revenue levels can occur. It is hard to imagine that such a contract form would not cause suspicion or could be justified in a court. Therefore, the transfer schedules will be restricted to being non-decreasing. Secondly, since there may be revenue levels that the producer believes to be impossible there is need for a lower bound to the transfer in any given state. If not, the distributor could make the payment infinitely negative in any state the producer believed to be impossible. If one of these states actually are possible under full awareness, the distributor would gain an infinite surplus in expectation. A reasonable level of this lower bound is zero as negative values imply that the producer would be paying the distributor despite already given away the good. This reasoning leads to the following assumption:

Assumption 2. *Weakly monotone and non-negative contracts: For all* π, $t(\pi) \geq 0$ *and* $t(\pi'') \geq t(\pi')$ *whenever* $\pi'' > \pi'$.

When chosen efforts are e and the transfer schedule is $t(\pi)$, the expected utility functions for the respective players are given by

$$U_D = E_S\pi(e,s) - E_S t(\pi(e,s)) - c(e_D), \tag{2.1}$$

$$U_P = E_{S|m} t(\pi(e,s)) - c(e_P). \tag{2.2}$$

Henceforth, the notation $E_X f(.)$ will refer to the expectation of function $f(.)$ when the expectation is taken over a state space X. Here that means that $E_S\pi(.)$ is the expectation of the revenue function over state space S. Note that the notation $S|m$ is the state space the producer perceives after receiving message m. Next, there are some assumptions made regarding how the state affects revenues and marginal impact of efforts. It turns out to be convenient to make a distinction between what constitutes a good or a bad state. Therefore, the following assumption relates to how a state affects the marginal productivity of effort.

Assumption 3. *If* $\pi(e,s') > \pi(e,s'')$, $s', s'' \in S$, *then* $\frac{\partial \pi(e,s')}{\partial e_i} > \frac{\partial \pi(e,s'')}{\partial e_i}$, $\forall\, i \in \{P, D\}$

The assumption states that if a state s' leads to a higher revenue than another state s'' for given efforts, then the marginal effect from effort is higher in state s' than in state s''. A possible justification is that if demand is high, there are higher rewards to be made from increasing the quality of the product than when demand is relatively low. Next, the following definition formalises the idea of the unknown contingency being 'good' or 'bad' news.

Definition 1.

(i) The contingency θ_2 *is said to be a positive contingency if*

$$\pi(e, (\tilde{\theta}_1, 1)) > \pi(e, (\tilde{\theta}_1, 0)),\ \forall\ e,\ \forall\ \tilde{\theta}_1.$$

(ii) The contingency θ_2 *is said to be a negative contingency if*

$$\pi(e, (\tilde{\theta}_1, 1)) < \pi(e, (\tilde{\theta}_1, 0)),\ \forall\ e,\ \forall\ \tilde{\theta}_1.$$

If the unknown contingency is a positive contingency, revenue would always be higher under its realisation than if it does not realise, and a negative contingency is defined in the converse way. By Assumption 3 it is also learned that if a positive contingency is realised, the marginal effect of efforts must be higher than if the contingency does not realise.

2.3 Observable Effort

First, consider the case where the effort choices of both players are observable and verifiable. The problem the distributor faces is to select a revelation decision, effort levels for both players and a transfer schedule that maximises his utility subject to the contract being accepted by the producer. Since efforts here are observable, effort levels can be specified in the contract. The producer will accept the contract if her utility from doing so is at least as large as the exogenously determined outside option, $\bar{U} \geq 0$, which leads to the individual rationality constraint (IR) stated below. The outlined problem is represented as follows:

$$\max_{m,t(\pi),e_D,e_P} E_S\pi(e,s) - E_S t(\pi(e,s)) - c(e_D), \tag{2.3}$$
$$s.t.\ (IR): \quad E_{S|m}t(\pi(e,s)) - c(e_P) \geq \bar{U}.$$

Recall that $E_{S|m}t(\pi(e,s))$ is the expected transfer when the expectation is taken according to the producer's post-communication awareness.

In order to find whether the distributor would like to reveal or not, I first analyse the optimal contracts and corresponding utility outcomes for the case where the distributor reveals the contingency. Subsequently, I analyse the outcomes when the contingency is not revealed and compare the two cases to find the distributor's optimal revelation decision.

2.3.1 Observable Effort and Revelation

Given revelation, I proceed in two steps to find equilibrium efforts and transfer schedule. By standard arguments, the IR constraint will in the optimal contract hold with equality, as otherwise the distributor could lower the transfer for some revenue level and be better off while still satisfying the constraint. So optimal efforts to implement can be found by solving (2.3) subject to the IR constraint holding with equality. After finding the optimal efforts, a transfer schedule can be found that minimises expected payment by the distributor while satisfying the IR constraint with equality. If such a contract exists, it must be optimal, since it induces optimal effort levels and satisfies the IR constraint with equality. An optimal contract here refers to one that maximises the distributor's utility subject to satisfying the individual rationality constraint.

Some useful notation is as follows. Denote the optimal efforts under revelation as $e^* = (e^*_D, e^*_P)$ and denote the expected joint surplus under revelation as $\psi^* = E_S\pi(e^*, s) - c(e^*_D) - c(e^*_P)$. Similarly, the efforts to be implemented under no revelation will be denoted by $e^{**} = (e^{**}_D, e^{**}_P)$. I will write ψ^{**} to be the expected surplus under efforts e^{**}, but where the expectation is over the full state space S, such that $\psi^{**} = E_S\pi(e^{**}, s) - c(e^{**}_D) - c(e^{**}_P)$.

As argued, optimal efforts are found by maximising the objective function while letting the IR constraint hold with equality. Substituting the binding IR constraint into (2.3) yields the unconstrained maximisation problem

$$\max_{e_P, e_D} E_S\pi(e, s) - \bar{U} - c(e_D) - c(e_P).$$

Taking a first order condition with respect to efforts gives

$$E_S \frac{\partial \pi(e, s)}{\partial e_i}\Big|_{e=e^*} = c'(e^*_i) \ \forall \ i \in \{P, D\}. \tag{2.4}$$

These are the optimal effort choices for the distributor to implement, conditional on participation by the producer. The reduced problem for the distributor after fixing efforts at e^* is

$$\max_{t(\pi)} E_S\pi(e^*, s) - E_S t(\pi(e^*, s)) - c(e^*_D),$$
$$s.t.: \ E_S t(\pi(e^*, s)) - c(e^*_P) \geq \bar{U}.$$

Since both players are risk neutral and efforts have been fixed, the remaining relevant feature of the contract is the expected payment. Furthermore, both players here have full awareness and so are taking the same expectation over the transfer schedule. The only constraint is that this expectation is equal to the producer's outside option and effort cost. Therefore, the optimal solution can be obtained by any contract such that the IR constraint holds with equality and the contract is non-negative and non-decreasing to satisfy Assumption 2. The utility for the producer is trivially $\bar{U}$ since the IR constraint is binding. The distributor receives the residual expected surplus after the producer has been paid. Then, the expected utility for the distributor in the case of revelation can be written as

$$\psi^* - \bar{U}. \tag{2.5}$$

2.3.2 Observable Effort and No Revelation

Next, consider the case where the distributor does not reveal the unknown contingency. Under asymmetric awareness, there is an issue regarding the fact that the optimal effort from the distributor's perspective is different from the one perceived as optimal from the producer's perspective. For this reason, the distributor's actions will be restricted so that they appear optimal from the perspective of the producer. This avoids problems of the producer becoming suspicious when observing an effort input from the distributor that is different from the one conjectured by the producer. Note that the revelation decision itself is not required to seem optimal as the producer does not perceive any absence of communication.

Assumption 4. *If the contingency is not revealed, the contract offer must solve the distributor's problem when both firms perceive S^P to be the true state space.*

Assumption 4 summarises the above discussion by requiring that the distributor's contract offer under no revelation solves the optimisation problem under limited awareness for both players. Using an equivalent argument to (2.4), it must then be the case that the efforts obey

$$E_{S^P} \frac{\partial \pi(e,s)}{\partial e_i}\Big|_{e=e^{**}} = c'(e_i^{**}),\ \forall\ i \in \{P, D\}. \tag{2.6}$$

Next, as the IR constraint in equilibrium always holds with equality under full awareness, Assumption 4 also implies that the IR constraint must appear to be satisfied with equality from the perspective of the producer. Fixing efforts at e^{**} and imposing that the IR constraint must hold with equality, the problem for the distributor to solve is

$$\max_{t(\pi)} E_S \pi(e^{**}, s) - E_S t(\pi(e^{**}, s)) - c(e_D^{**}), \tag{2.7}$$

$$s.t.:\ E_{S^P} t(\pi(e^{**}, s)) = c(e_P^{**}) + \bar{U}. \tag{2.8}$$

Recall from the revelation section that once efforts are fixed, there are only a discrete number of revenue levels possible. The distributor is aware of four possible realisations, whereas the producer is only aware of two of these. So when finding an optimal transfer schedule, the distributor only needs to consider the expectation over the transfer at the two revenue levels that the producer is aware of. For the remaining two revenue levels that the producer is unaware of, the distributor is maximising his utility by minimising the payment in these states subject to not violating any of the constraints. In other words, the distributor is solving

$$\min_{t(\pi)} E_S t(\pi(e^{**}, s)), \tag{2.9}$$
$$s.t.: \; E_{S^P} t(\pi(e^{**}, s)) = c(e_P^{**}) + \bar{U}.$$

Note the importance of the monotonicity constraint in setting the transfer schedule. The expected payment over the revenue realisations that the producer is aware of is tied down by the IR constraint (2.8), so the distributor cannot 'exploit' the producer over these transfer realisations. But he can more freely choose transfer values over revenue levels the producer does not believe to be possible, subject to satisfying the non-negativity and monotonicity constraints. Which the relevant constraints are depend on the impact of the unknown contingency were it to occur. Equilibrium contracts will be discussed for the cases of the unknown contingency being 'large' or 'small'.

Definition 2.

(i) A positive contingency is large if, for all e,

$$\pi(e, 11) > \pi(e, 01) > \pi(e, 10) > \pi(e, 00),$$

and it is small if for all e

$$\pi(e, 11) > \pi(e, 10) > \pi(e, 01) > \pi(e, 00).$$

(ii) A negative contingency is large if, for all e,

$$\pi(e, 10) > \pi(e, 00) > \pi(e, 11) > \pi(e, 01),$$

and it is small if for all e

$$\pi(e, 10) > \pi(e, 11) > \pi(e, 00) > \pi(e, 01).$$

A large positive contingency implies that the revenue levels that are unknown to the producer are both strictly larger than the revenue levels that are known to her. Crucially, this means that if the contingency occurs, the resulting revenue will necessarily be higher than the producer would have considered possible. Similarly, the realisation of a large negative contingency implies that revenue will necessarily be lower than the producer would have imagined. For small contingencies, there is not a strict ordering between the two revenue levels possible with an occurring unknown contingency and the two revenue levels possible when the contingency does not occur. It turns out that the shape of the optimal contract offer will depend on whether contingencies are large or small. Therefore, the analysis below derives in turn the optimal contracts for the various types of contingencies that have been defined.

2.3.2.1 Case 1: Large positive contingency

In this case, the unknown event would have such a large positive effect on revenues that if it occurs, revenues are always larger than they could have been had it not occurred. This means that the two revenue levels that the producer is unaware of are higher than the two levels the producer is aware of. Recalling the minimisation problem (2.9), the distributor wants to minimise payment in these two higher states subject to the IR and monotonicity constraints being satisfied. Because of the monotonicity constraint, this happens when the payments in the two known states are equalised at $c(e_P^{**}) + \bar{U}$ so that the transfer in the two higher unknown states can also be set at this level. If the payments in the known states were not equalised, by monotonicity the payments in the higher states would have to be higher than $c(e_P^{**}) + \bar{U}$. It must therefore be the case that an optimal contract, for any revenue level π, is

$$t(\pi) = c(e_P^{**}) + \bar{U}.$$

This means that the producer's payment expectation coincides with the payment the distributor expects to make, $c(e_P^{**}) + \bar{U}$. Using (2.1) with this transfer schedule and efforts e^{**} yields a distributor utility of

$$U_D = \psi^{**} - \bar{U}. \tag{2.10}$$

2.3.2.2 Case 2: Small positive contingency

Here the effect of the occurrence of the unknown contingency is still positive but lower than the effect of the known contingency. The first step towards finding an optimal contract is then to realise that in any optimal contract it must be the case that

$$t(\pi(e^{**}, 00)) = t(\pi(e^{**}, 01)) \tag{2.11}$$

and

$$t(\pi(e^{**}, 10)) = t(\pi(e^{**}, 11)). \tag{2.12}$$

This is because given any payment in the states known to the producer, this is the lowest possible payments in the unknown states that can be made given the monotonicity constraint, just as argued above regarding a large positive contingency. To find the implied expected transfer, note that the transfer in states (00) and (01) is paid out whenever any of these states occur, which happens with probability $(1 - p_1)$. Similarly, the transfer in states (10) and (11) is paid out with probability p_1. Therefore, the expected payment can be stated as

$$(1 - p_1)t(\pi(e^{**}, 00)) + p_1 t(\pi(e^{**}, 10)).$$

This is exactly the expectation over the transfer that the producer makes, $E_{S^P} t(\pi(e^{**}, s))$. According to the IR constraint (2.8), this value must be equal to the producer's cost plus outside option. Therefore, any contract that satisfies the IR with equality, satisfies the monotonicity constraint and satisfies $t(\pi(e^{**}, 00)) = t(\pi(e^{**}, 01))$ and $t(\pi(e^{**}, 10)) = t(\pi(e^{**}, 11))$ is an optimal contract.

To find the utility for the distributor, note that from (2.11) and (2.12), the expected payment made will be

$$E_S t(\pi(e^{**}, s)) = (1 - p_1)t(\pi(e^{**}, 00)) + p_1 t(\pi(e^{**}, 10)),$$

which simplifies to $c(e_P^{**}) + \bar{U}$ by the binding IR constraint. The distributor therefore pays in expectation the same amount that the producer believes to be paid. Plugging this expected payment into the distributor's objective (2.1) yields

$$U_D = \psi^{**} - \bar{U}. \tag{2.13}$$

2.3.2.3 Case 3: Negative contingency

Finally, consider the case where the effect of the occurrence of the unknown contingency has a negative effect on revenues. If the negative contingency is large, it must be that

$$t(\pi(e^{**}, 01)) = t(\pi(e^{**}, 11)) = 0,$$

since both revenue levels the producer is unaware of are lower than the levels she is aware of. Therefore, the transfers over these levels can be set to the lower bound 0 without constraining the transfers over the known revenue levels.

Consider next a small negative contingency. By an identical argument, $t(\pi(e^{**}, 01))$ can be set to 0 without tightening any constraints. Similarly, the transfer in the other state unknown to the producer, (11), will also be set as low as possible without violating monotonicity, which implies $t(\pi(e^{**}, 11)) = t(\pi(e^{**}, 00))$. Therefore, it becomes optimal for the producer to set the transfer in state (00) as low as possible and instead satisfy participation through payment in the highest state (10), as this lowers the payment that has to be made in state (11). So the shape of the optimal contract in this case must be

$$t(\pi) = \begin{cases} 0 & \text{if } \pi < \pi(e^{**}, 10), \\ \frac{1}{p_1}(c(e_P^{**}) + \bar{U}) & \text{if } \pi \geq \pi(e^{**}, 10), \end{cases}$$

where the payment in the highest state is chosen to just satisfy the IR constraint. Again, to find the resulting utility for the distributor, consider the expected payment that he makes. Since a payment is only made in the state (10) and this state occurs with probability $p_1(1 - p_2)$, the expected payment can be written as

$$\begin{aligned} E_S t(\pi(e, s)) &= p_1(1 - p_2) t(\pi(e^{**}, 10)), \\ &= p_1(1 - p_2) \frac{1}{p_1}(c(e_P^{**}) + \bar{U}), \\ &= (1 - p_2)(c(e_P^{**}) + \bar{U}). \end{aligned} \tag{2.14}$$

This expected payment is strictly less than the payment the producer expects to receive, which was not the case with a positive contingency. Previously, the monotonicity constraint prevented the distributor from exploiting the fact that he is aware of more possible revenue levels than the producer is. However, with a negative contingency, the revenue levels that the producer is unaware of are both lower than the highest possible revenue level they are jointly aware of. The distributor can then shift all payment needed to ensure participation to that highest state, so the monotonicity constraint will not constrain the distributor over the revenue levels the producer is unaware of. This means that the distributor's and producer's beliefs over payment are no longer aligned and the distributor can pay less in expectation to satisfy the IR constraint.

Plugging the expected payment into the distributor's objective (2.1) gives

$$\psi^{**} - \bar{U} + p_2(c(e_P^{**}) + \bar{U}). \tag{2.15}$$

2.3.3 Revelation Decision with Observable Effort

The equilibrium contracts for the various cases of contingencies under observable effort have now been derived. By comparing the resulting utilities for the distributor conditional on his revelation decision, it is possible to state conditions under which he will communicate the second contingency.

Proposition 1. *Consider the case of observable effort and suppose Assumptions 1-4 are satisfied.*

(i) The contingency is revealed whenever it is a positive contingency.

(ii) When the contingency is a negative contingency, it is revealed whenever the following condition holds:

$$p_2(\bar{U} + c(e_P^{**})) < \psi^* - \psi^{**}. \tag{2.16}$$

Furthermore, there exist parameter values for which the contingency is revealed and there exist parameter values for which the contingency is not revealed.

Proof. (i) Under revelation, from (2.5) the distributor's utility is $\psi^* - \bar{U}$. Under no revelation, from (2.10) and (2.13), whenever there is a positive contingency the distributor's resulting utility is $\psi^{**} - \bar{U}$. Next note that since e_P^* and e_D^* are by definition the efforts that maximise $E_S\pi(e,s) - c(e_P) - \bar{U} - c(e_D)$, it must be that $\psi^* > \psi^{**}$. Hence, it follows that the distributor will choose to reveal the contingency.

(ii) With a negative contingency, the distributor's utility under revelation is still $\psi^* - \bar{U}$ from (2.5), while under no revelation the distributor's utility is $\psi^{**} - \bar{U} + p_2(c(e_P^{**}) + \bar{U})$ from (2.15). Comparing these two resulting utilities leads to condition (2.16). The last part of the statement will be proven by constructing two examples demonstrating that both revelation and no revelation are possible equilibrium choices.

Suppose that $\bar{U} = 0$ and $E_S\pi(e,s) = 0, \forall\ e$, but that $E_{S^P}\pi(e,s) > 0, \forall\ e$, and $e^{**} > 0$. This means that it is a valueless product ex ante regardless of effort inputs, although the producer believes it to be valuable. Trivially, $e^* = (0,0)$ and therefore the distributor's utility under revelation will be 0. The distributor's utility under no revelation will be

$$E_S\pi(e^{**}, s) - c(e_D^{**}) - (1 - p_2)c(e_P^{**}) - (1 - p_2)\bar{U},$$

which simplifies since $E_S\pi(e^{**}, s) = 0$ and $\bar{U} = 0$:

$$-c(e_D^{**}) - (1 - p_2)c(e_P^{**}) < 0.$$

The distributor will therefore choose to reveal the unknown contingency.

In the next example suppose that $e^* = e^{**} > 0$. Recall from Assumption 1, part (*iv*), that $\pi(e, s') \neq \pi(e, s''), \forall\ s', s'' \in S$. Since $e^* = e^{**}$, it follows that $\phi^* = \phi^{**}$. Distributor utility from revelation is then $\phi^* - \bar{U}$ while utility from no revelation is $\phi^* - \bar{U} + p_2(c(e_P^{**}))$. By inspection, the no revelation utility is higher and the distributor will not reveal the contingency. □

Whenever there is good news to be told, the distributor chooses to communicate this to the producer. The reason is that by revealing the information, the distributor is able to offer a contract that implements the optimal effort levels. If the information is not revealed, the distributor must according to Assumption 4 offer a contract that seems optimal to the producer, which means that the effort levels the contract aims to achieve may be different from the optimal effort levels. Furthermore, whenever there is a positive contingency, the distributor is not able to exploit the fact that he has superior awareness. This is because the producer's and the distributor's expectations over payments in this case coincide under any optimal contract.

However, when there is bad news to be told it could be advantageous for the distributor to conceal the contingency. This is because now the expectations of the two players over payments are no longer coinciding. If the contingency is not revealed, the distributor pays less than the producer believes to be paid. There is therefore a trade-off between implementing optimal efforts and payment needed to ensure participation. Looking at the key condition (2.16), the lefthand side represents the expected exploitation benefit since if the second contingency is realised the distributor avoids the expected payment $c(e_P^{**}) + \bar{U}$. The righthand side represents the efficiency benefits from having optimal efforts. Note that the distributor becomes the full residual claimant to the efficiency benefits since he is the contract proposer and holds the producer to her perceived reservation utility.

The proposition shows that either effect could dominate, depending on model parameters. A high probability of the unknown contingency occurring increases the exploitation benefit and therefore makes revelation less likely. A higher outside option for the producer also increases the exploitation benefit, since by not revealing the distributor can avoid a larger expected transfer. The sensitivity of optimal efforts with respect to the unknown contingency is also important as this determines the magnitude of the efficiency loss of not revealing the contingency. These points are well illustrated by the examples provided in the proof of part (*ii*). If the product is valueless, but the producer believes it to have value, the contingency must be revealed. This is because by revealing, the outcome must be zero for both players since the product is valueless. If not revealing, the distributor has to pay the producer despite no revenues being raised. The second example shows that if efforts are inelastic to the unknown contingency, the contingency will not be revealed. This is because there are then no efficiency benefits from revelation but there may be exploitation benefits.

2.4 Unobservable Effort

Consider now the case where producer effort is unobservable.[2] The main difference between this case and the observable effort case is that the distributor can no longer impose an effort level on the producer. Instead, the producer will now optimally choose an effort level conditional on the signed contract. Therefore, the distributor must take into account the producer's endogenous effort choice when designing the contract. This adds an incentive compatibility (IC) constraint to the distributor's problem. This constraint requires that the producer effort level that the distributor wishes to implement indeed is the one the producer wishes to choose, i.e.

$$e_P \ \in \arg\max_{e_P} E_{S|m} t(\pi(e, s)) - c(e_P). \tag{2.17}$$

As previously, the revelation decision will be found by solving the model for revelation and no revelation and comparing the resulting utilities.

2 Note that distributor effort is kept observable. The set-up is ill-suited to unobservable effort of the contract proposer because the distributor could then adjust his own effort such that the revenue levels the producer believes to be possible could in fact never occur.

2.4.1 Unobservable Effort and Revelation

The distributor's contracting problem conditional on revealing the unknown contingency and effort being unobservable is:

$$\max_{t(\pi),e_D,e_P} \; E_S\pi(e,s) - E_S t(\pi(e,s)) - c(e_D), \tag{2.18}$$

$$s.t.\ (IR):\ \ E_S t(\pi(e,s)) - c(e_P) \geq \bar{U}, \tag{2.19}$$

$$s.t.\ (IC):\ \ e_P \in \arg\max_{e_P} E_S t(\pi(e,s)) - c(e_P). \tag{2.20}$$

First, consider the effort choice that the distributor wants to induce. If he is able to induce the same unconstrained optimal effort that was induced with observable effort while satisfying the IR constraint with equality and satisfying the IC constraint, then this must be the optimal effort to induce with unobservable effort. Assume therefore that e^* is the effort choice that the distributor seeks and if a transfer schedule can be found that satisfies the constraints then this must be an optimal contract. The problem for the distributor is then

$$\max_{t(\pi)} E_S\pi(e^*,s) - E_S t(\pi(e^*,s)) - c(e_D^*),$$

$$s.t.:\ \ E_S t(\pi(e^*,s)) - c(e_P^*) \geq \bar{U},$$

$$s.t.:\ \ e_P^* \in \arg\max_{e_P} E_S t(\pi(e_P,e_D^*,s)) - c(e_P).$$

I claim that there exists a contract of the form $t(\pi) = \pi + k \ \forall \ \pi$ that solves this problem. Since the proposed contract is continuous and differentiable, transform the IC constraint by taking a first order condition of the objective in (2.17):

$$E_S \frac{\partial t(\pi(e^*,s))}{\partial \pi} \frac{\partial \pi(e,s)}{\partial e_P}\bigg|_{e=e^*} = c'(e_P^*).$$

If $\partial t(\pi(e^*,s))/\partial\pi = 1 \ \forall \ \pi$ is imposed, the optimal effort condition seen in (2.6) is mimicked and therefore optimal effort is induced. A schedule that has this property is indeed $t(\pi) = \pi + k$, where k is a constant. Next, the contract must satisfy the IR constraint with equality:

$$E_S t(\pi(e^*,s)) = \bar{U} + c(e_P^*),$$

$$E_S\pi(e^*,s) + k = \bar{U} + c(e_P^*),$$

$$\therefore k = \bar{U} + c(e_P^*) - E_S\pi(e^*,s).$$

This can be substituted into the contract form to give

$$
\begin{aligned}
t(\pi(e^*, s)) &= \pi(e^*, s) + k, \\
&= \pi(e^*, s) + \bar{U} + c(e_P^*) - E_S\pi(e^*, s), \\
&= \pi(e^*, s) - E_S\pi(e^*, s) + \bar{U} + c(e_P^*). \qquad (2.21)
\end{aligned}
$$

It follows that the 'selling the project' contract is an optimal contract. The producer gets the full amount of the generated revenue but has to pay the ex ante expected revenue level to the distributor. This gives the producer incentives to choose the optimal effort level as she fully internalises the effects of her effort, while the contract still allows the distributor to reap his maximum possible expected share of the surplus.

This is not the unique optimal contract, there is indeed a number of possible optimal contracts. What these contracts must have in common is that there must be a strict increase in the transfer schedule at some point, and the size of this increase must be such that the producer chooses the optimal effort. Nevertheless, it is clear that it is possible to implement the optimal effort while satisfying the constraints, so the resulting utility for the distributor will be

$$\psi^* - \bar{U}. \qquad (2.22)$$

2.4.2 Unobservable Effort and No Revelation

If the distributor does not reveal the unknown contingency, the contracting problem for the distributor is

$$
\begin{aligned}
&\max_{t(\pi), e_D, e_P} E_S\pi(e, s) - E_S t(\pi(e, s)) - c(e_D), \\
&s.t.: \ E_{S^P} t(\pi(e, s)) - c(e_P) \geq \bar{U}, \\
&s.t.: \ e_P \in \arg\max_{e_P} E_{S^P} t(\pi(e, s)) - c(e_P).
\end{aligned}
$$

As previously, assume that the optimal efforts are implementable which can be verified later. Also, recall that by Assumption 4 the optimal efforts here are the ones as perceived by the producer, which still is e^{**} as given by (2.6). Then the problem for the distributor is

$$
\begin{aligned}
&\max_{t(\pi)} E_S\pi(e^{**}, s) - E_S t(\pi(e^{**}, s)) - c(e_D^{**}), \\
&s.t.: \ E_{S^P} t(\pi(e^{**}, s)) - c(e_P^{**}) \geq \bar{U}, \\
&s.t.: \ e_P^{**} \in \arg\max_{e_P} E_{S^P} t(\pi(e_P, e_D^{**}, s)) - c(e_P).
\end{aligned}
$$

It turns out that similar to the observable case, the optimal contract will depend on whether the contingency is large or small, in the sense of Definition 2. Each relevant case will now be covered in turn.

2.4.2.1 Case 1: Large positive contingency

Similar to the reasoning with observable effort, the distributor needs to select transfers over the revenue levels that the producer is aware of to satisfy the IR and the IC constraints. Here both these revenue levels are lower than the two possible revenue levels that the producer is unaware of.

The shape of the transfer schedule that relaxes the IC constraint the most is a step function form, as e.g. displayed in Figure 2.2, because this ensures that the producer does not have incentives to make a local effort deviation. A local upward deviation of effort would increase effort cost with no change in expected payment, while a downward local deviation would significantly lower payment for low savings in effort cost. Therefore, the step function form ensures there to be only two potential effort deviations to consider. First, the producer could deviate downwards such that she only receives payment in the state (10) but instead saves on effort cost. The best effort for the producer to choose in this case, denoted by $e_P^{\downarrow}$, is the lowest possible effort such that the lower transfer level is ensured to occur in the state (10). Therefore, $e_P^{\downarrow}$ is defined by

$$\pi(e_D^{**}, e_P^{\downarrow}, 10) = \pi(e^{**}, 00). \tag{2.23}$$

The other candidate deviation is an upward deviation such that the producer receives the payment $t(\pi(e^{**}, 10))$ with certainty, but instead has to pay a higher effort cost. This optimal deviation effort $e_P^{\uparrow}$ is the lowest possible effort such that the higher transfer is perceived to be guaranteed and so is defined by

$$\pi(e_D^{**}, e_P^{\uparrow}, 00) = \pi(e^{**}, 10). \tag{2.24}$$

Recall that the best solution for the distributor when the unknown contingency is positive is a flat schedule where $t(\pi) = c(e_P^{**}) + \bar{U},\ \forall\ \pi \geq \pi(e^{**}, 00)$. This contract is displayed in Figure 2.1 where it is visualised that the flat contract minimises the payment in the higher states that are unknown to the producer.[3] With the flat

3 Since neither player believes revenue levels below $\pi(e^{**}, 00)$ to be possible, the transfer could equivalently be set to any value for these revenue levels, as long as it does not affect the monotonicity constraint.

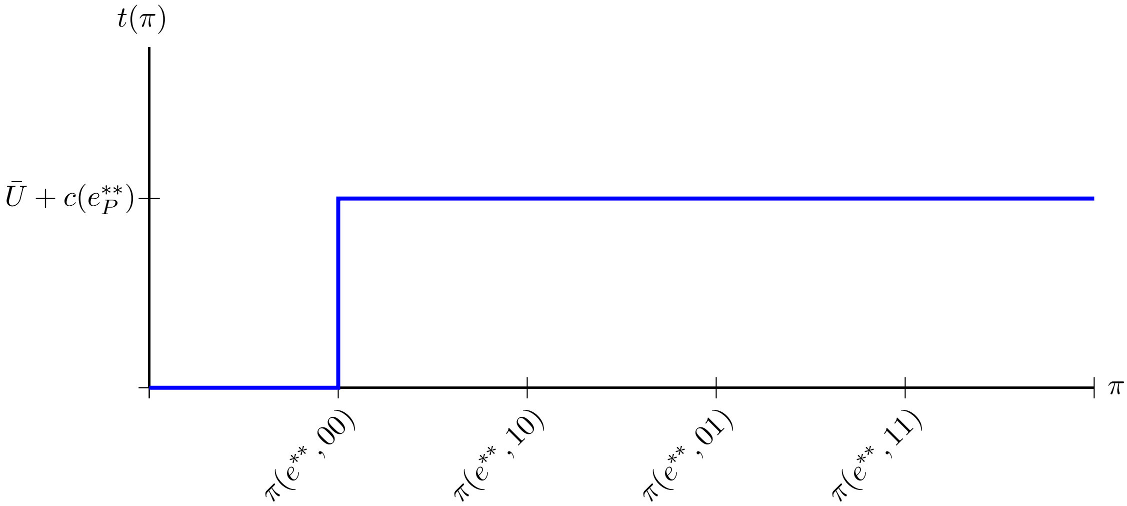

Figure 2.1: Large positive contingency and non-binding IC constraint

contract, an upward effort deviation would never be profitable for the producer as this would not increase transfers paid. Therefore, there is only need to consider a potential downward deviation from this contract. As argued, the optimal potential deviation would be to $e_P^{\downarrow}$ as defined in (2.23). Note that the distributor would always set $t(\pi) = 0$ for $\pi < \pi(e^*, 00)$, both to relax constraints and minimise payments, so a downward deviation will mean that no transfer is paid in the state (00).

The IC constraint implies that the utility for the producer from choosing effort e_P^{**} must be weakly larger than the utility of choosing effort $e_P^{\downarrow}$:

$$\begin{aligned}
U_P(e_P^{**}) &\geq U_P(e_P^{\downarrow}), && \\
E_{S^P} t(\pi(e^{**}, s)) - c(e_D^{**}) &\geq E_{S^P} t(\pi(e_D^{**}, e_P^{\downarrow}, s)) - c(e_P^{\downarrow}), && \text{by (2.1), (2.2)} \\
\bar{U} &\geq E_{S^P} t(\pi(e_D^{**}, e_P^{\downarrow}, s)) - c(e_P^{\downarrow}), && \text{(binding IR)} \\
\bar{U} &\geq p_1 t(\pi(e_D^{**}, e_P^{\downarrow}, 10)) - c(e_P^{\downarrow}), && (t(\pi(e_D^{**}, e_P^{\downarrow}, 00)) = 0) \\
\bar{U} &\geq p_1 t(\pi(e^{**}, 00)) - c(e_P^{\downarrow}), && \text{by (2.23)}
\end{aligned}$$

$$\therefore \ \frac{\bar{U} + c(e_P^{\downarrow})}{t(\pi(e^{**}, 00))} \geq p_1. \tag{2.25}$$

In particular, if $p_1 > [\bar{U} + c(e_P^{\downarrow})]/[\bar{U} + c(e_P^{**})]$, the flat transfer schedule preferred by the distributor will not be incentive compatible. The optimal contract will therefore be a function of p_1; if p_1 is high the distributor will lower payment in the state (00) while increasing payment in (10) to simultaneously satisfy the IR and IC constraints. By the monotonicity constraint, this necessarily increases payment in the states (01)

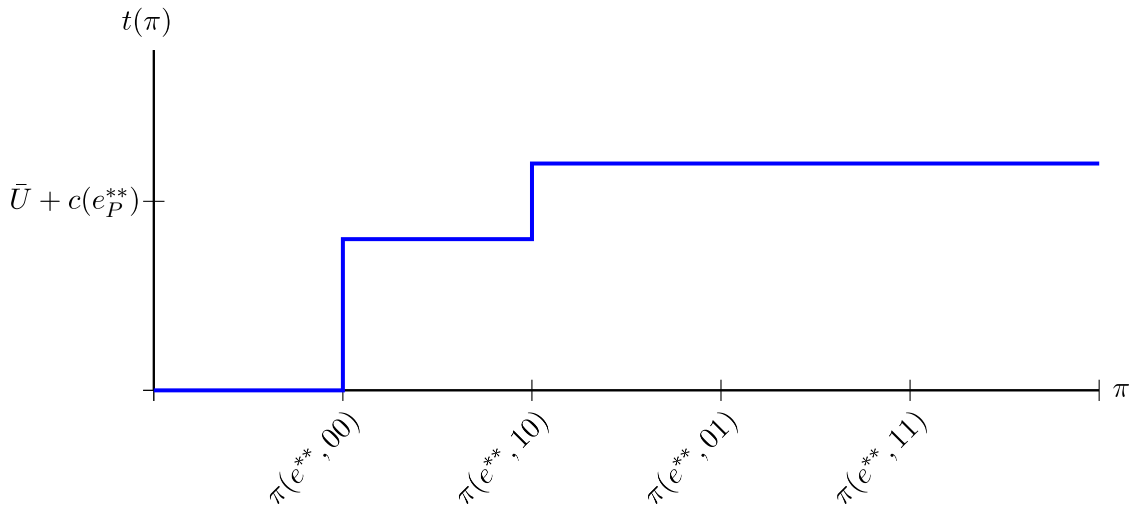

Figure 2.2: Large positive contingency and binding IC constraint

and (11). The shape of the optimal contract in this case is displayed in Figure 2.2. Note how the monotonicity constraint forces the payments in the unknown states to be higher than with a flat transfer schedule. If p_1 is low, a flat schedule as in Figure 2.1 is feasible and is the preferred schedule by the distributor.

In the case where the IC constraint binds, the distributor still wants to keep payment in the state (10) as low as possible, so will always choose $t(\pi(e^{**}, 00))$ such that (2.25) holds with equality. The binding IC and IR constraints imply that the transfers in the constrained case must be

$$t(\pi(e^{**},00)) = \frac{1}{p_1}[\bar{U} + c(e_P^{\downarrow})], \tag{2.26}$$

$$t(\pi(e^{**},10)) = \frac{1}{p_1}\left[\bar{U} + c(e_P^{**}) - \frac{1-p_1}{p_1}[\bar{U} + c(e_P^{\downarrow})]\right]. \tag{2.27}$$

According to the discussion, the optimal contract can then be characterised as follows:

If $p_1 \leq \frac{\bar{U}+c(e_P^{\downarrow})}{\bar{U}+c(e_P^{**})}$,

$$t(\pi) = \begin{cases} 0 & \text{if } \pi < \pi(e^{**},00), \\ c(e_P^{**}) + \bar{U} & \text{if } \pi \geq \pi(e^{**},00). \end{cases}$$

If $p_1 > \frac{\bar{U}+c(e_P^{\downarrow})}{\bar{U}+c(e_P^{**})}$,

$$t(\pi) = \begin{cases} 0 & \text{if } \pi < \pi(e^{**},00), \\ \frac{1}{p_1}(c(e_P^{\downarrow}) + \bar{U}) & \text{if } \pi \in (\pi(e^{**},00), \pi(e^{**},10)), \\ \frac{1}{p_1}\left[\bar{U} + c(e_P^{**}) - \frac{1-p_1}{p_1}[\bar{U} + c(e_P^{\downarrow})]\right] & \text{if } \pi \geq \pi(e^{**},10). \end{cases}$$

Now turn to the utility implications of this contract. With a flat contract, the expected payment will be $c(e_P^{**}) + \bar{U}$ which implies that the utility for the distributor is $\psi^{**} - \bar{U}$ as in the previous cases. If p_1 is large enough such that the IC constraint is binding, the payment will be state dependent as described above. By noting that the lower payment occurs with probability $(1-p_1)(1-p_2)$ and the higher payment with complementary probability, it can be worked out that the expected payment is

$$E_S t(\pi(e,s)) = (1 - p_2 + \frac{p_2}{p_1})[\bar{U} + c(e_P^{**})] - \frac{(1-p_1)p_2}{p_1^2}[\bar{U} + c(e_P^{\downarrow})]. \tag{2.28}$$

The utility for the distributor will then be

$$U_D = \psi^{**} - \bar{U} + p_2 \frac{1-p_1}{p_1}\left[\frac{1-p_1}{p_1}\bar{U} - c(e_P^{**}) + \frac{1}{p_1}c(e_P^{\downarrow})\right], \tag{2.29}$$

which is strictly less than $\psi^{**} - \bar{U}$ whenever $p_1 > [\bar{U} + c(e_P^{\downarrow})]/[\bar{U} + c(e_P^{**})]$ is true.

2.4.2.2 Case 2: Small positive contingency

Recall from the observable effort section that with a small positive contingency, any contract that satisfies the IR and monotonicity constraints must be optimal, as they imply the same payoff for the distributor. With unobservable effort some of these contracts may not be incentive compatible, but such a contract must always exist. For example, the 'selling the project' contract is always incentive compatible. So, there is always a contract payoff-equivalent to the ones in the observable effort case that can be implemented.

This means that the utility for the distributor will be the same as in the observable effort case,

$$U_D = \psi^{**} - \bar{U}. \tag{2.30}$$

2.4.2.3 Case 3: Negative contingency

With a negative unknown contingency and unobservable effort, in contrast to observable effort, it now matters whether there is a small or large negative contingency. When there is a large negative contingency, the non-negativity constraint on transfers is trivially the relevant constraint and the transfer at the unknown levels can be set to 0. The transfer over the known revenue levels can be set in any way that satisfies the IR constraint with equality and satisfies the IC constraint. A solution to that problem always exists, for example through the 'selling the project' contract.

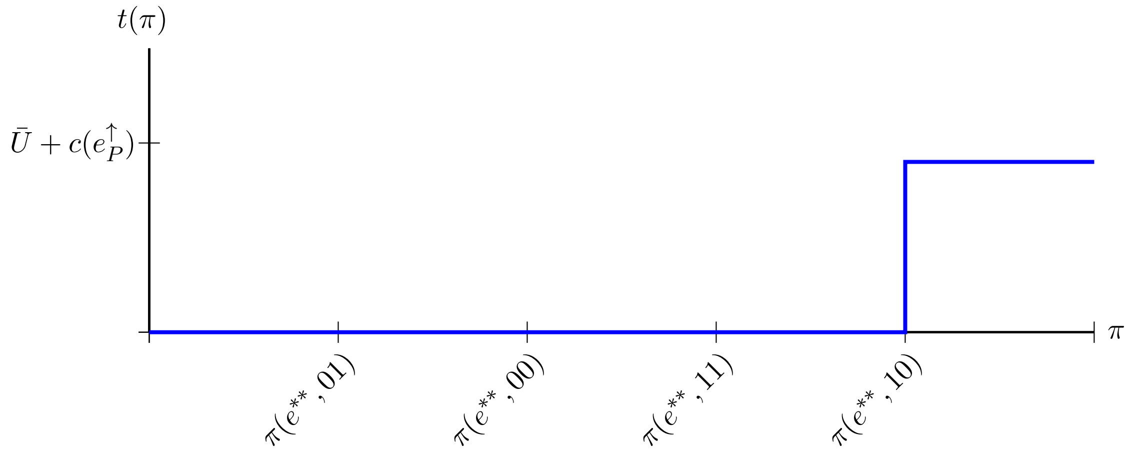

Figure 2.3: Negative contingency and non-binding IC constraint

If instead there is a small negative contingency, the reasoning from the observable effort case reveals that the distributor would prefer to satisfy the IR constraint through payment in the state (10) rather than in (00) as this lowers payment in the unknown state (01). However, the inclusion of the IC constraint may make this contract infeasible. Consider the contract that only pays in the state (10). Firstly, it is clear that a downward effort deviation would never be profitable as it means that a transfer would never be paid. However, an upward deviation to $e_P^{\uparrow}$ as defined by (2.24) may be profitable, as this would give the producer a perceived guaranteed payment of $t(\pi(e^{**}, 10))$. The relevant constraint is therefore ensuring that the producer does not exert too much effort.

First, note that the IR constraint binding means that the transfer in the high state must satisfy

$$\frac{1}{p_1}\left[\bar{U} + c(e_P^{**}) - (1-p_1)t(\pi(e^{**}, 00))\right]. \tag{2.31}$$

The IC constraint is then found by ensuring that the producer's perceived utility from choosing e_P^{**} is weakly larger than the utility of choosing $e_P^{\uparrow}$:

$$\begin{aligned} U_P(e_P^{**}) &\geq U_P(e_P^{\uparrow}), \\ E_{S^P} t(\pi(e^{**}, s)) - c(e_P^{**}) &\geq t(\pi(e^{**}, 10)) - c(e_P^{\uparrow}), && \text{by (2.1), (2.2)} \\ \bar{U} &\geq t(\pi(e^{**}, 10)) - c(e_P^{\uparrow}), && \text{(binding IR)} \\ \bar{U} &\geq \frac{1}{p_1}\left[\bar{U} + c(e_P^{**}) - (1-p_1)t(\pi(e^{**}, 00))\right] - c(e_P^{\uparrow}), && \text{by (2.31)} \end{aligned}$$

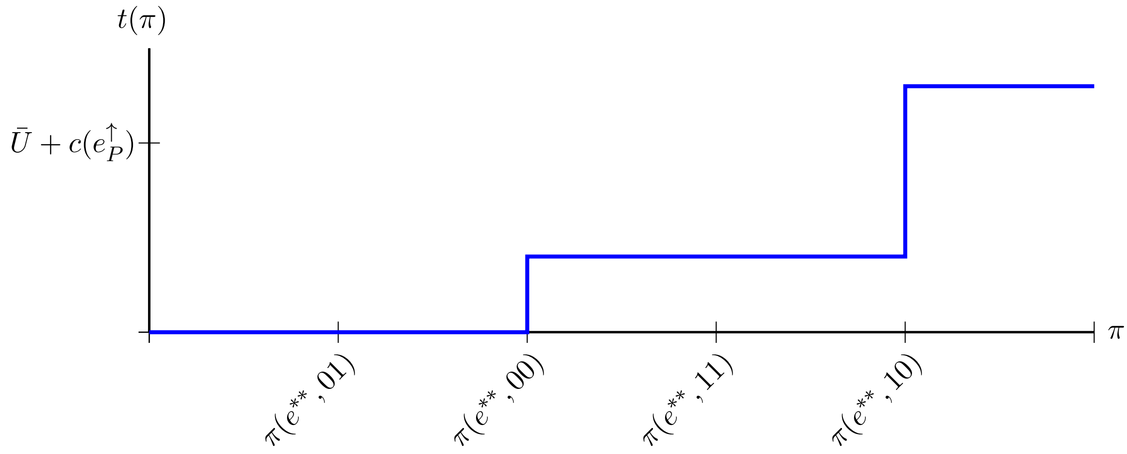

Figure 2.4: Negative contingency and binding IC constraint

$$\therefore\ p_1 \geq \frac{\bar{U} + c(e_P^{**}) - t(\pi(e^{**}, 00))}{\bar{U} + c(e_P^{\uparrow}) - t(\pi(e^{**}, 00))}. \tag{2.32}$$

Now p_1 needs to be high enough in order for the contract to be incentive compatible. If the high revenue state is likely to occur, there is less to gain for the producer from increasing effort to always achieve this level. In this case, the distributor can offer his preferred unconstrained contract which is displayed in Figure 2.3. However, if the high revenue state is unlikely, there is more to gain from always achieving this very high payment and so the producer has an incentive to deviate upwards. In this instance, the distributor is forced to shift some payment to the state (00) to limit the incentives for the producer to deviate. The transfer schedule then looks as displayed in Figure 2.4. Since the distributor prefers the payment in state (00) to be as low as possible, when $t(\pi(e^{**}, 00)) = 0$ is not feasible he would choose a transfer $t(\pi(e^{**}, 00))$ such that the IC constraint (2.32) is satisfied with equality. Rearranging the IC constraint gives

$$t(\pi(e^{**}, 00)) = \bar{U} + \frac{1}{1 - p_1} c(e_P^{**}) - \frac{p_1}{1 - p_1} c(e_P^{\uparrow}). \tag{2.33}$$

Using the binding IR constraint to find what this implies for the transfer in state (10), the optimal contracts can be characterised as:

If $p_1 \geq \frac{\bar{U} + c(e_P^{**})}{\bar{U} + c(e_P^{\uparrow})}$,

$$t(\pi) = \begin{cases} 0 & \text{if } \pi < \pi(e^{**}, 10), \\ \frac{1}{p_1}(c(e_P^{**}) + \bar{U}) & \text{if } \pi \geq \pi(e^{**}, 10). \end{cases}$$

If $p_1 < \frac{\bar{U}+c(e_P^{**})}{\bar{U}+c(e_P^{\uparrow})}$,

$$t(\pi) = \begin{cases} 0 & \text{if } \pi < \pi(e^{**}, 00), \\ \bar{U} + \frac{1}{1-p_1}c(e_P^{**}) - \frac{p_1}{1-p_1}c(e_P^{\uparrow}) & \text{if } \pi \in (\pi(e^{**}, 00), \pi(e^{**}, 10)), \\ \bar{U} + c(e_P^{\uparrow}) & \text{if } \pi \geq \pi(e^{**}, 10). \end{cases}$$

Noting that the lower payment occurs with probability $1 - p_1p_2$ and the higher payment occurs with probability p_1p_2, it can be shown that the expected payment is

$$(1 - p_2)\left[\bar{U} + c(e_P^{**})) + p_1p_2[\bar{U} + \frac{1}{1 - p_1}c(e_P^{**}) - \frac{p_1}{1 - p_1}c(e_P^{\uparrow})]\right]. \tag{2.34}$$

The utility of the distributor can then be shown to be

$$\psi^{**} - \bar{U} + p_2\left[(1 - p_1)\bar{U} + \frac{1 - 2p_1}{1 - p_1}c(e_P^{**}) + \frac{p_1^2}{1 - p_1}c(e_P^{\uparrow})\right]. \tag{2.35}$$

2.4.3 Revelation Decision with Unobservable Effort

The optimal contracts under unobservable effort have now been solved for and the resulting utilities for the distributor can be used to find the revelation decision. To simplify exposition, Table 1 summarises the utilities for the distributor under unobservable effort and no revelation. Recall that under revelation, the utility for the distributor is always $\psi^* - \bar{U}$ as in (2.22).

Proposition 2. *Consider the contracting model with unobservable effort and suppose Assumptions 1-4 are satisfied.*

(i) The contingency is revealed whenever it is a positive contingency.

(ii) When there is a negative contingency, there exist parameter values for which the contingency is revealed and there exist parameter values for which the contingency is not revealed.

(iii) If the contingency is revealed under observable effort, it must also be revealed under unobservable effort. Furthermore, if there is a small negative contingency and $p_1 < [\bar{U} + c(e_P^{**})]/[\bar{U} + c(e_P^{\uparrow})]$*, there exist cases for which the contingency is revealed under unobservable effort but not under observable effort.*

Table 2.1: Distributor utility under unobservable effort and no revelation

Contingency case	U_D **when not revealing**
Large positive & $p_1 \leq \frac{\bar{U}+c(e_P^{\downarrow})}{\bar{U}+c(e_P^{**})}$	$\psi^{**} - \bar{U}$
Large positive & $p_1 > \frac{\bar{U}+c(e_P^{\downarrow})}{\bar{U}+c(e_P^{**})}$	$\psi^{**} - \bar{U} + p_2 \frac{1-p_1}{p_1} \left[\frac{1-p_1}{p_1} \bar{U} - c(e_P^{**}) + \frac{1}{p_1} c(e_P^{\downarrow}) \right]$
Small positive	$\psi^{**} - \bar{U}$
Small negative & $p_1 \geq \frac{\bar{U}+c(e_P^{**})}{\bar{U}+c(e_P^{\uparrow})}$	$\psi^{**} - \bar{U} + p_2(c(e_P^{**}) + \bar{U})$
Small negative & $p_1 < \frac{\bar{U}+c(e_P^{**})}{\bar{U}+c(e_P^{\uparrow})}$	$\psi^{**} - \bar{U} + p_2 \left[(1-p_1)\bar{U} + \frac{1-2p_1}{1-p_1} c(e_P^{**}) + \frac{p_1^2}{1-p_1} c(e_P^{\uparrow}) \right]$
Large negative	$\psi^{**} - \bar{U} + p_2(c(e_P^{**}) + \bar{U})$

Proof. (i) Consider a positive contingency and note that the utility for the distributor when revealing is the same regardless of effort observability, $\psi^* - \bar{U}$. Next, note that whenever there is a positive contingency, the utility for the distributor when not revealing the contingency is the same, or lower, when effort is unobservable compared to when effort is observable. This follows from the discussion and from that the set of feasible contracts is in the unobservable case a subset of the feasible contracts in the observable case. Since the contingency was always revealed under observable effort by Proposition 1, it therefore follows that the contingency will always be revealed under unobservable effort as well.

(ii) Note that if $p_1 \geq [\bar{U} + c(e_P^{**})]/[\bar{U} + c(e_P^{\uparrow})]$, the utility for the distributor when not revealing is the same under unobservable and observable effort. By Proposition 1, both revelation and no revelation was possible with a negative contingency and the proof did not make use of any assumptions on p_1. Therefore either revelation or no revelation could also occur under unobservable effort.

(iii) First note that from (i) and (ii) it has been established that for a positive contingency the revelation decisions under observable and unobservable effort coincide. For a negative contingency, consider the condition for revelation under observable effort

$$\psi^* - \psi^{**} = p_2(c(e_P^{**}) + \bar{U}), \tag{2.36}$$

versus the condition under unobservable effort

$$\psi^* - \psi^{**} > p_2 \left[(1-p_1)\bar{U} + \frac{1-2p_1}{1-p_1} c(e_P^{**}) + \frac{p_1^2}{1-p_1} c(e_P^{\uparrow}) \right]. \tag{2.37}$$

It has been established that, when $p_1 < [\bar{U} + c(e_P^{**})]/[\bar{U} + c(e_P^{\uparrow})]$, the expression for the righthand side in (2.37) is strictly smaller than (2.36). Furthermore, by choosing an appropriate $\pi(.)$ the lefthand side of the equations can be freely chosen. Hence, there must exist parameter values such that the following set of inequalities hold:

$$p_2 \left[(1 - p_1)\bar{U} + \frac{1 - 2p_1}{1 - p_1} c(e_P^{**}) + \frac{p_1^2}{1 - p_1} c(e_P^{\uparrow}) \right] < \psi^* - \psi^{**} < p_2(c(e_P^{**}) + \bar{U}),$$

which implies that the contingency is revealed under unobservable effort, but not under observable effort. □

The results stated in the proposition rest on the result that the contract under the two cases of observability will only differ when the IC constraint is binding. For good news, a binding IC constraint makes no difference to the distributor's utility under revelation as there are numerous incentive compatible contracts that yield the observable outcome. But under no revelation and a positive contingency, a binding IC constraint makes the distributor worse off with unobservable effort than observable effort since it implies that the contract must be upward sloping. This leaves higher payments to the producer in the states that the producer is unaware of. Since the contingency was always revealed with observable effort it must therefore also always be revealed under unobservable effort.

When there is a negative contingency and the IC constraint is not binding, it similarly is the case that the contracts under unobservable effort and observable effort are identical, which means that both revelation and no revelation can occur. However, when the IC constraint is binding the contracts are differing such that under no revelation the distributor is worse off with unobservable effort than with observable effort. This demonstrates that the distributor's revelation decision may depend on whether effort is observable. If the distributor chooses to not reveal the contingency when effort is observable he may then choose to reveal the contingency if effort was unobservable. This because he becomes less able to exploit his superior awareness when designing the contract.

It is interesting to note that unobservability of effort may constrain the distributor whenever the contingency is not revealed. Looking from the perspective of a standard principal-agent model, this result is different to the result from the standard model where observability of effort does not matter for the principal when both players are risk neutral. In the current setting with asymmetric awareness, observability does matter for the principal/distributor since unobservability constrains the set of feasible contracts that can be offered. In particular, when there is bad news to be shared the concern is that effort incentives are too strong under the contract the distributor would like to offer, while under good news the effort incentives are too weak under the preferred contract.

2.5 Welfare Analysis

The discussion so far has concerned model predictions. This section analyses the welfare implications of these predictions. I will focus on comparing welfare outcomes from the two different assumptions on effort observability discussed. Welfare will here be taken to mean joint surplus for the players.

First, note that from the perspective of the supply chain, the decision that maximises total surplus is to always reveal the information. This is because if information is revealed, both players have a correct expectation of the revenue prospects of the product and the distributor can, without causing suspicion, offer a contract that implements the jointly optimal efforts. If a negative contingency is not revealed, both players exert too much effort into the product. If instead a positive contingency is not revealed, both players exert too little effort. Therefore, if wishing to find out which case of observability that is most efficient, it suffices to consider under which case it is most likely that information is revealed.

Part (*iii*) of Proposition 2 provides the crucial information. It shows that if a contingency is revealed under observable effort, it must always be revealed under unobservable effort. However, the reverse is not true, as there exist parameter values for which information is revealed under unobservable effort but not under observable effort. Therefore, the set of parameter values that yields outcomes where information is being revealed when effort is observable is a strict subset of the corresponding set of parameter values for unobservable effort. In this sense, it is more likely that information is revealed under unobservable effort than under observable effort. It follows that unobservable effort is more efficient for the supply chain than observable effort.

Next, consider each firm's preference regarding observability of effort. The only cases in which the observability assumption alters the revelation decision is for those parameter values where the contingency is revealed under unobservable effort but not revealed under observable effort. If the distributor reveals the contingency, he obtains identical outcomes in either case of observability as in both cases the jointly optimal efforts are implemented and the producer is paid the reservation payment. Since this revelation outcome was feasible in both cases, but the distributor preferred not to reveal in the observable case, it must be the case that he is weakly better off when effort is observable. Regarding the producer, note that she would always prefer the contingency to be revealed, since in that case she receives her reservation payment in expectation, while if the contingency is not revealed she may receive less. Then, since it has been established that information is communicated more often in the unobservable case, the producer must prefer effort to be unobservable.

These welfare results are summarised in Proposition 3. Proof is omitted as the results follow from the preceding discussion.

Proposition 3. *Suppose that Assumptions 1-4 are satisfied.*

(i) The joint surplus for the players is weakly higher under unobservable than under observable effort.

(ii) The distributor's surplus is weakly higher under observable effort than under unobservable effort. The producer's surplus is weakly higher under unobservable effort than under observable effort.

The first result highlights the tension between efficiency and exploitation that the distributor faces in the revelation decision. By revealing the information, the total surplus is increased, which also benefits the distributor. If information is not revealed, total surplus is lower, which hurts the distributor. But there may now also be an exploitation benefit where the distributor does not have to pay the producer's full outside option. The analysis has shown that the distributor is less able to avoid payments to the producer when effort is unobservable, which is why information is more often revealed in this case. Therefore, unobservable effort is revenue enhancing for the supply chain as it promotes communication. Regarding the players' individual preferences over observability, there is a tension in the sense that the producer prefers effort be unobservable, while the distributor prefers it to be observable.

The proposition discusses the welfare impacts for the firms in the supply chain. If considering social welfare, one may want to take into account the effect on other stakeholders, in particular consumers. The impact on consumers will depend on how efforts and transfer schedules affect the downstream market. The effort level that is jointly optimal for the supply chain need not be the optimal effort level from a consumer perspective. However, if effort levels do not alter the marginal cost of producing the product, consumers ought to benefit from higher effort as higher effort may not translate into higher price through marginal cost effects. If this is the case, consumers would always prefer outcomes that imply higher effort levels. In particular, this would mean that consumers prefer the contingency to be revealed if it is a positive contingency, as this increases efforts, but that negative contingencies are not revealed. As such, consumers would prefer effort to be observable, as this leads to less communication of negative contingencies. However, sharper statements about the impact on consumers would require adding more structure to the downstream market.

2.6 Conclusion

This chapter has presented a model of asymmetric awareness in a contractual relationship where the producer of a good is unaware of the possibility of an event occurring that would affect the expected profitability of the product. It has been shown that whenever the impact of this event would be positive on the outcome of the project if occurring, this information will be shared by the distributor. This is also the jointly optimal outcome for the supply chain. If instead there is bad news to be told, the information may not be shared, which leads to overinvestment in efforts. The reason for this lack of information sharing lies in that the superiorly informed distributor is then able to design a contract which pays less to the producer in expectation than the producer expects to be paid.

By comparing the outcomes of the model under different assumptions regarding the observability of the producer's effort, it is shown that there will be more communication when effort is unobservable, as unobservable effort constrains the distributor's ability to exploit his superior awareness. This leads to the welfare result that the joint surplus of the firms is higher with unobservable effort. The results presented hinge on the assumption that the distributor possesses full bargaining power. This means that an optimal contract for the distributor will implement the first best efforts and there is no hold-up problem that could lead to inefficiently low effort inputs. It also depends on the agents being risk neutral, which means that there is no efficiency loss from effort being unobservable. If the producer was risk averse, it could potentially become jointly optimal not to share information as that could move the level of efforts closer to the jointly optimal level.

The setting of unawareness in the model implies that there is no signalling aspect of the contractual shape. One motivation for this is that the unaware producer does not know what superior information the distributor could have, even if knowing that the distributor has superior information, so there would be no basis on which to make inference. Possibilities for exploring signalling under unawareness could be to use methods developed in decision theory by Karni and Vierø (2017), where an unaware agent has an attitude towards the unknown as a primitive. This may be an interesting avenue for future research to explore, as an alternative to traditional signalling models. This kind of analysis could also be used to approach the issue of repeated interaction and reputational concerns. If a distributor has in the past been involved in interactions where suppressed information could be suspected, the producer may attempt to make inference from any lack of information sharing or from the offered contractual shape. However, it is not clear in the current set-up how this inference would be formed, given the producer's unawareness. The current framework ought to be most robust in situations where repeated interaction is unlikely or where it is difficult ex post to deduce whether information has been suppressed.

3 Markets with Differentiated Add-on Products

3.1 Introduction

In some markets, the products on offer are not sold exclusively in one single version and are offered alongside upgrade or add-on options. The premium versions of the products might be offered at the same time as the basic versions, or at a subsequent opportunity. Examples of upgrades and add-ons include purchasing add-on insurance when renting a car, utilising a credit card overdraft facility or using the mini-bar when renting a hotel room. More broadly, one may consider products that are offered to consumers with different quality levels; when buying a car, for example, consumers can purchase the basic version of the car and are also offered a menu of upgrade options, such as leather seating or a turbo engine. I will henceforth make the distinction between a *core product* and an *add-on product*. The key feature is that the add-on product must not be purchased in isolation; it can only provide consumers with utility if purchased in combination with the core product.

The relationship between add-on prices and their marginal costs can vary a lot between markets, making it important to understand the strategy of price setting with add-on products. For example, hotel mini-bar prices tend to be far higher than marginal costs for providing this service, while newspapers sometimes include special weekend magazines free of charge to their subscribers. While in many markets add-on products are included free of charge, much regulatory and academic attention has been focused on add-on products priced far higher than their cost levels. The concern is consumers being ripped off by high prices, as consumers may not be fully aware of the consequences when committing to purchases, for example regarding bank overdraft charges or mutual funds fees.[1] However, these types of concerns are challenged by the Chicago-style argument that high add-on prices ought to be matched by low core product prices as firms compete to attract consumers to their profitable add-ons. This leaves consumers and firms on average as well off as in a competitive single product

1 See discussion e.g. in Armstrong and Vickers (2012) and Stango and Zinman (2014) on bank overdraft charges, and Khorana *et al.* (2009) and Malkiel (2013) on mutual fund and asset management fees.

market, only leading to distributive concerns.[2]

In this chapter, I consider a novel justification for firms' pricing power in add-on products; the possibility that add-on products are differentiated independently of the core product differentiation. To illustrate, when buying travel insurance, consumers may after having selected which provider to purchase from, be presented with an opportunity for additional coverage. The offered additional coverage may not be identical across providers or even be correlated with the core product content. One provider may offer high additional medical coverage and another provider could offer additional legal coverage. In fact, general insurance add-ons were the subject of an FCA investigation[3] due to the suggestion that consumers were insufficiently informed about add-on terms and that competition was not working properly. The investigation was published in 2015 and resulted in a ban on opt-out selling of add-ons as well as measures requiring firms to improve the information supplied to consumers about add-ons at the point of core product purchase. Other relevant examples of differentiated add-ons might include computer software, which comes in basic options but also in premium versions and with subsequent upgrades, where the components of the premium content or upgrades may differ across brands depending on their relative strengths.[4]

I will in this paper investigate the impact on market outcomes of add-on products being differentiated across firms. As such, I construct a model of a price-setting duopoly market where firms are supplying differentiated core and add-on products. Since add-on product markets tending to catch regulatory and academic attention are characterised by limited consumer information, my model incorporates this feature by allowing a fraction of consumers to be uninformed about add-on preferences and prices. I then investigate the relationship between add-on prices and the level of information. The set-up of differentiated add-ons also allows me to develop this model further, as the add-on differentiation gives rise to a natural justification for rational consumer information acquisition. If add-ons are differentiated, there is value to consumers from investigating which add-on product they prefer, which provide them with incentives to acquire information. In section 3.3, an exogenously given fraction

2 This type of argument is formalised by Lal and Matutes (1994) in their model of loss-leader pricing.
3 Financial Conduct Authority (2014).
4 Bonatti (2011) mentions the case of Central Processing Units, where Intel has comparative strength in clock speed while AMD has comparative strength in cache memory. Consumers may then have different tastes for quality across firms depending on which characteristics that are important to that consumer.

of consumers is informed, while in section 3.4 this reasoning is exploited to allow consumers to rationally decide whether to become informed or not.

I find that the differentiation in add-on prices leads to above-cost pricing of add-ons, even in the case of a fully informed consumer population. These inefficiently high add-on prices are socially harmful as they prevent efficient matching between firms and consumers. High add-on prices imply that firms offer lower add-on surplus, leading consumers to underweight the surplus created when they purchase the add-on product. Consumers therefore too often purchase their preferred core product without purchasing an add-on. There is therefore inefficient matching between consumers and firms, even if firms and equilibrium prices are symmetric. Furthermore, add-on prices are increasing in the fraction of consumers that are uninformed. An important corollary is that matching is worsened when there are more uninformed consumers and not simply because uninformed consumers may inefficiently select their firm. With more uninformed consumers, higher add-on prices make the informed consumers more likely to make inefficient firm choices, creating externalities between the consumer groups. I also point out that in the setting with differentiated add-on products, the profit irrelevance result of add-ons does not apply. Firms do compete away part of the add-on profit in the core product market, but incompletely so as long as there are sufficient amounts of informed consumers. This is due to the add-on differentiation leading to a firm's inframarginal consumers being more profitable than its marginal consumers, which softens competition.

In section 3.4, I show that when the sample of informed consumers is determined by rational information acquisition, the insight that add-on prices fall as more consumers become informed of add-on information may no longer hold. The mechanism behind this rests on the idea that the consumers with most to gain from becoming informed are those with weakest firm preference. Because of this, firms have incentives to compete fiercely for these consumers as they are disproportionally likely to be marginal. As more consumers become informed the average informed consumer has a stronger firm preference, lowering incentives to compete for these consumers and add-on prices may rise. I also show that from a social welfare perspective, there is underinvestment in information acquisition as consumers do not internalise the effect on profits from them acquiring information. Conversely, from a consumer welfare perspective, there is excessive investment in information acquisition arising from the fact that consumers do not internalise the externality exerted on other consumers through their impact on equilibrium prices.

The remaining part of this section outlines the relevant literature. Section 3.2 presents the general set-up of the model used in this chapter. Section 3.3 analyses the model under the assumption that the fraction of consumers that is informed is randomly assigned. Section 3.4 presents an analysis of the case where consumers rationally decide whether to become informed. Section 3.5 concludes.

Literature

In many instances of add-on products, consumers may be imperfectly informed about the add-on product at the time of deciding which firm to purchase from. This has led the literature to conclude that add-ons in such markets are expensively priced because consumers are locked in to a firm after having bought that firm's core product, as shown by Verboven (1999).[5] He shows in a duopoly model that having consumers uninformed of add-on prices leads to monopoly level add-on prices and finds that this model suits his data better than a model with fully informed consumers.

A result in Verboven's paper is that highly priced add-ons are irrelevant from the firms' perspective since any profit made on the add-on products is being competed away on the core products which are priced low, potentially below cost. This type of irrelevance result was earlier shown in Lal and Matutes (1994) in their analysis of loss-leader pricing. Ellison (2005) departs from this irrelevance result, by explaining in a model of add-on pricing that abnormal profits can be explained by an adverse selection effect where a price-cutting firm is disproportionally likely to attract 'cheapskates' that do not purchase the add-on product. This argument rests on correlation between firm and add-on preferences, as it is assumed that consumers with a strong firm preference are also more likely to purchase the add-on product. The set-up of Ellison's model is similar to mine but crucially assumes positive correlation between core preference and preference over add-ons, which drives his results. The add-on profit irrelevance result has also been shown not to apply in some settings with boundedly rational consumers, as by Gabaix and Laibson (2006). In their model firms make profits on add-ons as firms that attempt to undercut on prices re-

5 In fact, in Verboven's model, consumers are locked in to a firm simply by visiting a firm due to the presence of a search cost, similar to the Diamond Paradox introduced by Diamond (1971).

ceive disproportionally many 'cheapskates' that do not buy add-ons. Other examples of papers where naive consumers are ripped off are Armstrong and Vickers (2012) who discusses contingent charges in financial services and Grubb (2015) who analyses bill shock regulation in a model with inattentive consumers. Heidhues *et al.* (2017) shows how a price-floor for core products can lead to profit-relevant and endogenously unobserved add-ons.

My model is also related to the literature on upgrade pricing, following Johnson and Myatt (2003) and Johnson and Myatt (2006). They analyse quality increments as upgrades using the demand profile approach as introduced in Wilson (1993). This convenient approach is not applicable in my model due to the differentiation of add-on products. Johnson and Myatt find that firms set above-cost add-on prices even under competition, a result arising from that they consider that firms compete in quantities rather than in prices.

The model is related to the literature on competitive non-linear pricing, where firms offer multiple quality levels at different prices.[6] The baseline monopoly version was presented by Mussa and Rosen (1978) and this basic model has been expanded to oligopoly models, as by Armstrong and Vickers (2001) and Rochet and Stole (2002). Relevant for my purposes, both of these last papers present a non-discrimination theorem saying that when brand preferences and quality valuations are independent, firms would in competitive equilibrium offer all quality levels at the same markup over costs. The implication for add-on pricing is that under these conditions, the full information outcome ought to be that add-ons are priced at marginal costs. This is also the outcome in the full information game in Verboven (1999). The paper in the competitive non-linear pricing literature most related to my work is Bonatti (2011), who presents a model where consumers may have brand-specific tastes for quality, which is similar to the assumption of differentiated add-ons that I maintain. While his main focus is on which endogenous quality schedule that are offered by firms in equilibrium, it is shown that stronger brand-specific tastes for quality lead to higher marginal prices at all quality levels.

Finally, Armstrong (2015) surveys models in which there are informed consumers and uninformed consumers and discusses what type of externalities these consumers

6 In this literature, quality can interchangeably be interpreted as quantity and a common focus is on comparing non-linear pricing with the benchmark of linear pricing. However, in my add-on set-up, it does not make much sense to think of 'linear' pricing as a potential policy intervention.

exert on each other. He presents a model of add-on pricing where all consumers have the same linear demand for each firm's add-on product. This leads to a number of differences to my model with differentiated add-ons. For example, unlike in my model, the equilibrium add-on prices are independent of the core product differentiation. In Armstrong's version firms only make profits from the presence of add-ons if consumers have wary beliefs (to be discussed later), while in my model firms make profits also when consumers have passive beliefs.

3.2 Model

I analyse a duopoly market where each firm (1 and 2) supplies a core product as well as an add-on product. Both products are produced by both firms at zero marginal cost.[7] Firms compete by setting prices p for the core product and p_U for the add-on product.[8] Throughout the analysis, it will be supposed that the second firm follows the (yet to be found) equilibrium strategies p^* and p_U^*, while firm 1's deviation incentives from these strategies will be considered to derive the equilibrium. Note that attention is restricted to pure strategies and symmetric equilibria. It is assumed that a consumer cannot purchase an add-on product without already having purchased the same firm's core product, so it is not possible to combine an add-on product with a different firm's core product. I also make the one-stop shopping assumption that a consumer is not allowed to purchase both firms' core products.

There is a unit mass continuum of risk neutral and expected utility maximising consumers uniformly distributed on a Hotelling line of length 1, where $\theta \in [0, 1]$ denotes a consumer's position on the line. The firms are located at each of the ends of the line. There is a differentiation parameter $t > 0$ which is the standard Hotelling transport cost. Consumers are also differentiated with respect to their preferences over the add-on products. A consumer either values the add-on at 0 or values it at the value $w > 0$. There are therefore four add-on types of consumers: $\{HH, HL, LH, LL\}$. The first entry refers to value at firm 1 and the second entry refers to value at firm 2, where the value can be high (w) or low (0).[9] I will be making some symmetry restrictions on the distribution over add-on types. The probability of being an HL

7 Prices can therefore be thought of as markups over marginal costs.

8 Where U refers to upgrade, which is one potential interpretation of the add-on product.

9 So e.g. a consumer being of type HL would imply that this particular consumer values firm 1's add-on product at w and firm 2's add-on product at 0.

type is assumed equal to the probability of being an LH type, and the probability of being an HH type is assumed equal to the probability of being an LL type. This implies that the unconditional probability of having a high add-on value at a firm is a half and that the probability of being an HL type, denoted as ρ, will be sufficient to describe the distribution. This parameter represents the degree of differentiation over add-on products, as it relates to the probability of a consumer having a strict firm preference over add-on products. Note the assumption of independence of product preferences within a firm, so that a consumer's core preference does not convey any information about that consumer's add-on preference and vice versa.

A consumer can be either informed or uninformed of the add-on products. When a consumer is informed, the consumer knows the add-on prices that both firms have set and also the consumer's value for the add-on product at each firm. If the consumer is uninformed, she does not know either her values at each firm or the prices that both firms have set. Instead, the uninformed consumers make rational conjectures about what prices the firms have set and use the prior distribution over add-on values to form beliefs over the add-on value at each firm.

The utility for a consumer buying from firm 1 and 2 respectively is given by

$$\begin{aligned} U_1 &= v - p_1 - t\theta + \max\{w_1 - p_U, 0\}, \\ U_2 &= v - p^* - t(1-\theta) + \max\{w_2 - p_U^*, 0\}, \end{aligned} \tag{3.1}$$

where w_i represents the add-on value for the consumer at firm i and v is a valuation for the core products fixed across consumers that will throughout be assumed large enough to ensure a covered market.

The first move in the model is by the firms who simultaneously choose both their prices. Informed consumers then observe these choices, as well as their core and add-on product preferences, while uninformed consumers only observe the chosen core prices as well as their core product preferences. All consumers next decide which firm to purchase the core product from (the core market is by assumption covered). Uninformed consumers now discover the add-on information at their chosen firm and all consumers choose whether to purchase the add-on product.

3.3 Random Selection

In this section, there is an exogenously determined selection of consumers that is uninformed about the add-on products. A fraction $\sigma \in [0, 1]$ of consumers will be informed while the complementary fraction $(1 - \sigma)$ will be uninformed of the add-on products. The level of σ is commonly known across both consumers and firms. The draw of consumers that are informed is assumed to be independent of their preferences over the core and add-on products, so the fact that a consumer is informed does not convey any information regarding that consumer's add-on or core preferences. Recall from the model introduction that an informed consumer knows both prices and values of each add-on product, while an uninformed consumer knows neither but instead makes rational conjectures about prices and values.

An important issue in this setting regards how uninformed consumers are updating their beliefs when facing an unexpected core price from a firm. Since a core price different from the conjectured equilibrium does not occur on the equilibrium path, it is not immediate what beliefs consumers would have over a firm's add-on price given this observation. Among the set of possible belief rules, two prominent alternatives are *passive beliefs* and *wary beliefs*. Wary beliefs as introduced by McAfee and Schwartz (1994) in the context of vertical contracting assume that consumers believe that a firm will play a best response when setting its add-on price, given the deviation in the core price and the equilibrium profile. Passive beliefs instead assume that consumers do not alter their equilibrium conjectures about add-on prices when seeing unexpected core prices. Throughout this chapter, I will be assuming passive beliefs, since this aids exposition and the main results are robust to which of these belief systems that are used. Passive beliefs may be justified by consumer naivité since it is a stringent requirement for consumers to calculate a firm's best response given a deviation. Consumer naivité may be reasonable in the current setting as this could be the reason that some consumers are uninformed in the first place.

To be transparent about the impact of consumer beliefs, an analysis of the case of wary beliefs is provided in appendix section 3.A.5. Generally, the insight from this analysis is that add-on prices will be lower and core prices higher when consumers have wary beliefs than when they have passive beliefs. This is because consumers will interpret a lower core price as that the firm has set a higher add-on price. Therefore, deviations become less attractive and competition softened, leading to higher core prices which translate into lower add-on prices. Aside from this effect, the main qualitative results from the present section are maintained.

3.3.1 Solving the Model

The model is solved by considering the behaviour of the different consumer types in turn. When solving the model, attention will be restricted to interior cases in the sense that each firm always faces some demand from every add-on type. This restriction requires a large enough t such that the consumers with strongest core product preference (at the ends of the Hotelling line) always choose the firm they prefer the core product at.[10]

Beginning with uninformed consumers, these consumers will base their firm choices on the verifiable information on core prices and preferences that are available to them as well as on their conjectures about add-on prices. Since uninformed consumers with passive beliefs in a candidate symmetric equilibrium will always have symmetric beliefs about the firms' add-on prices and the values at each firm, their decision will only depend on the core information.

An uninformed consumer will purchase the core product from firm 1 if her expected utility from doing so is larger than the expected utility of purchasing from firm 2. Using (3.1), the following condition for an uninformed consumer to purchase from firm 1 is obtained:

$$v - p_1 - t\theta + E(\max\{w_1 - p_U^*, 0\}) > v - p^* - t(1-\theta) + E(\max\{w_2 - p_U^*, 0\}).$$

Since consumers expect add-on prices of p_U^* at both firms, the expectation here is with respect to the add-on preferences for each firm. Due to the symmetry of beliefs, the expected add-on surplus is equal at both firms and will drop out of the expression. Rearranging gives the condition

$$\theta < \frac{1}{2} + \frac{1}{2t}(p^* - p_1)$$

which specifies for what positions on the Hotelling line a consumer would purchase from firm 1. Next, recall that consumers are uniformly spread over the Hotelling line.

10 This assumption restricts attention to the most interesting case and avoids some equilibrium existence concerns. For instance, in the limit case where core products are homogenous, given $\sigma > 0$ there is no symmetric equilibrium in pure strategies. To see this, note that symmetric Bertrand competition over HH and LL types implies $p^* = p_U^* = 0$ and hence $\pi^* = 0$. But firm 1 then has a profitable deviation of setting a positive total price that would see some demand from HL types. Hence, there cannot exist a symmetric equilibrium.

Therefore, the probability of an uninformed consumer purchasing from firm 1 can be written as

$$\frac{1}{2} + \frac{1}{2t}(p^* - p_1). \tag{3.2}$$

This expression is valid as long as it is contained in $[0, 1]$, which is guaranteed under the maintained assumption that both firms always face demand from all add-on types of consumers.

Next, informed consumers take into account the offered add-on prices and the realised add-on value they have at each firm when they decide which firm to purchase the core product from. The choice that the consumers make will then depend on which add-on type they are. For illustration, suppose that a consumer is an HL type so that she has a high add-on value at firm 1 but a low add-on value at firm 2. This consumer will choose firm 1 if her utility from doing so is larger than the utility she would gain by buying at firm 2, conditional on being an HL type. Using the utility functions from (3.1), this is the case if the condition

$$v - p_1 - t\theta + w - p_U > v - p^* - t(1 - \theta)$$

holds, which can be rearranged to obtain the condition on θ for an informed HL type consumer to purchase from firm 1:

$$\theta < \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + w - p_U).$$

An equivalent derivation can be made for all consumer types, and using the fact that θ is uniformly distributed on $[0, 1]$, the following demands conditional on add-on type are obtained:

$$\begin{aligned} Pr(\text{Buy } 1|HH) &= \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + p_U^* - p_U), \\ Pr(\text{Buy } 1|HL) &= \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + w - p_U), \\ Pr(\text{Buy } 1|LH) &= \frac{1}{2} + \frac{1}{2t}(p^* - p_1 - w + p_U^*), \\ Pr(\text{Buy } 1|LL) &= \frac{1}{2} + \frac{1}{2t}(p^* - p_1). \end{aligned} \tag{3.3}$$

To find the profit firm 1 can expect to make off an uninformed consumer that has purchased its core product, note that the consumer will only purchase the add-on product if she is an HL or HH type, which happens with probability $\rho + (1/2 - \rho) =$

1/2. Therefore the expected profit for the firm per uninformed consumer that it attracts is

$$p_1 + \frac{1}{2}p_U.$$

For the informed consumers, the expected profitability is conditional on add-on type. The HH and HL consumers who purchase the core product from firm 1 will also purchase the add-on product from firm 1, whereas the LH and LL types will only buy firm 1's core product. Therefore, defining the events $\{I, N\}$ as a consumer being informed or not informed respectively, the profit function for firm 1 can be written as

$$\begin{aligned}\pi_1 = (1-\sigma)&\left[(p_1 + \frac{1}{2}p_U)Pr(1|N)\right]\\ &+ \sigma(p_1 + p_U)\left[(\frac{1}{2} - \rho)Pr(1|HH, I) + \rho Pr(1|HL, I)\right]\\ &+ \ \sigma p_1\left[\rho Pr(1|LH, I) + (\frac{1}{2} - \rho)Pr(1|LL, I)\right].\end{aligned}$$

Where $Pr(1|N)$ refers to the probability of purchasing the core product at firm 1 conditional on not being informed and the other expressions are accordingly defined. Using the expressions from (3.2) and (3.3) and summing up, the explicit expression for the profit function becomes

$$\begin{aligned}\pi_1 = (1-\sigma)&\left[(p_1 + \frac{1}{2}p_U)[\frac{1}{2} + \frac{1}{2t}(p^* - p_1)]\right]\\ &+ \sigma(p_1 + p_U)\left[\frac{1}{4} + \frac{1}{4t}(p^* - p_1 + 2\rho(w - p_U^*) + p_U^* - p_U)\right]\\ &+ \ \sigma p_1\left[\frac{1}{4} + \frac{1}{4t}(p^* - p_1 - 2\rho(w - p_U^*))\right].\end{aligned} \tag{3.4}$$

This profit function conveys some key intuition about firms' price setting incentives in this environment. When considering the demand for firm 1 arising from uninformed consumers, note that firms' add-on prices are not at all present. This is because of two reasons. First, uninformed consumers by construction do not observe add-on prices and so cannot base their choice of firm on the actual add-on prices. Second, the consumers will purchase the add-on product whenever $p_U \leq w$, which will always hold in equilibrium, so the amount of uninformed consumers purchasing the add-on product is entirely insensitive to add-on price. This means that changes in add-on prices have no impact on a firm's demand, suggesting that increasing the number of uninformed consumers ought to lead to higher add-on prices.

Next, consider the demand from informed consumers that are buying both the core product and the add-on product at firm 1:

$$\frac{1}{4} + \frac{1}{4t}(p^* - p_1 + 2\rho(w - p_U^*) + p_U^* - p_U), \tag{3.5}$$

which in equilibrium becomes

$$\frac{1}{4} + \frac{\rho}{2t}(w - p_U^*). \tag{3.6}$$

Similarly, the demand from informed consumers that are buying only the core product at firm 1 becomes in equilibrium

$$\frac{1}{4} - \frac{\rho}{2t}(w - p_U^*). \tag{3.7}$$

The firms split the core market equally between them so that each firm gets a measure of 1/2 consumers in any symmetric equilibrium. If the group of consumers that are purchasing the core product from firm 1 were a random selection of consumers, there ought to be a quarter demand for the core plus add-on offering at firm 1 and a quarter demand for the core only offering. However, looking at (3.6) and (3.7) it can be seen that whenever $p_U^* < w$ there is a shift in demand towards the add-on and core option compared to the demand with a random allocation of consumers. This is due to the fact that some consumers are selecting into a firm they have a high add-on value at and means that each firm faces a higher demand for their add-on product than the unconditional expectation would suggest.

3.3.2 Equilibrium Analysis

The derived profit function from the previous section is now used to analyse symmetric market equilibrium. Equilibrium is found by letting both firms maximise the profit function (3.4) by simultaneously choosing their core and add-on prices.

Proposition 1. *Suppose* $t > w/(1+4\rho)$.

(i) If $\sigma < 2t/(w+2t)$*, then in any symmetric equilibrium, equilibrium prices are given by*

$$p^* = t - \frac{1}{2}w,$$
$$p_U^* = w.$$

(ii) If $\sigma > 2t/(w+2t)$*, then in any symmetric equilibrium, equilibrium prices are given by*

$$p^* = t\frac{\sigma(2+4\rho)-1}{\sigma(1+4\rho)} - \frac{2\rho}{1+4\rho}w,$$
$$p_U^* = 2t\frac{1-\sigma}{\sigma(1+4\rho)} + \frac{4\rho}{1+4\rho}w.$$

Proof. Taking a first derivative of (3.4) with respect to p_1 and imposing the symmetric equilibrium requirements $p_1 = p^*, p_U = p_U^*$, one obtains

$$p^* = t - \frac{1}{2}p_U^*.$$

Similarly, taking a first derivative of (3.4) with respect to p_U and imposing symmetry gives

$$p_U^* = \frac{t\sigma^{-1} + 2\rho - p^*}{1+2\rho}.$$

Solving the set of simultaneous equations in p_U^* and p^* yields

$$p_U^* = 2t\frac{1-\sigma}{\sigma(1+4\rho)} + \frac{4\rho}{1+4\rho}w, \tag{3.8}$$
$$p^* = \frac{\sigma(1+4\rho)-1+\sigma}{\sigma(1+4\rho)}t - \frac{2\rho}{1+4\rho}w. \tag{3.9}$$

Whenever $\sigma > 2t/(w+2t)$, this is the unique stationary point and therefore unique equilibrium candidate.[11] When $\sigma < 2t/(w+2t)$, (3.8) yields $p_U^* > w$ and hence there is no stationary point such that $p_U^* < w$. Furthermore, due to no consumer valuing the add-on above w, the strategy $p_U^* = w$ strictly dominates $p_U^* > w$. Hence, the only candidate for equilibrium is $p_U^* = w$, in which case (3.8) and (3.9) already reveals that the best response in core prices will yield $p^* = t - (1/2)w$.

Finally, under (3.8) and (3.9), for symmetric prices, the profit function (3.4) applies whenever $t > w/(1+4\rho)$, since this guarantees that the add-on demands in (3.3) are strictly bounded by 0 and 1. □

11 Appendix section 3.A.1 demonstrates that the second order conditions are satisfied at these prices.

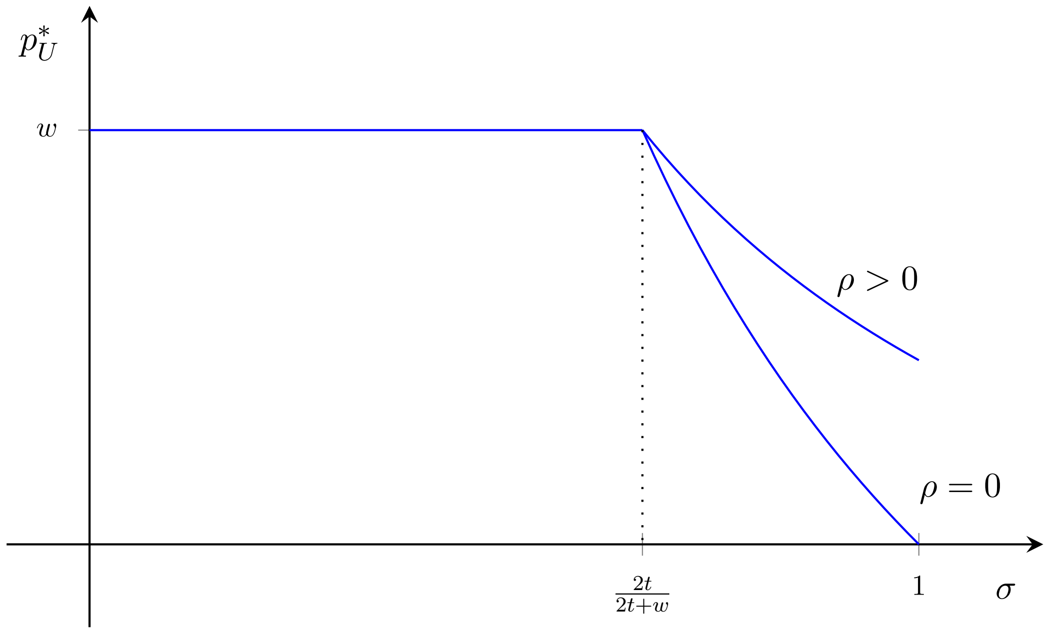

Figure 3.1: Equilibrium add-on price

The proposition states that for low levels of informed consumers, the firms find it optimal to charge the monopoly add-on price.[12] This is reasonable as with high levels of uninformed consumers, total demand is insensitive to changes in the add-on price and the firms optimally increase the add-on price as much as possible. However, as the level of informed consumers increases, the equilibrium add-on price starts to fall while the core price rises (it is easily shown that $(\partial p^*)/(\partial \sigma) > 0$ and $(\partial p_U^*)/(\partial \sigma) < 0$). This is illustrated in Figure 3.1. The reason is that with more informed consumers, demand becomes more sensitive to changes in the add-on prices. At a threshold level of σ, there will be an incentive to decrease the add-on price to attract more of the informed consumers, at the cost of lost profit on the uninformed consumers who do not react to changes in the add-on price. The add-on price then smoothly moves towards the level $(4\rho w)/(1+4\rho)$ which is reached when all consumers are informed of the add-on information. The market outcome will have marginal cost pricing of the add-on only in the case where all consumers are informed and the add-on product is homogenous.

A key result is that even in the case where all consumers are informed, the add-on price is held above marginal cost. The reason for this is that the add-on products

12 Note that equilibrium existence is not demonstrated. When t is close to $t > w/(1+4\rho)$, it is possible that firm 1 has a profitable deviation where it stops selling to LH types. However, a large enough value of t ensures that the proposed equilibrium exists.

are differentiated. If add-ons had been homogenous, the intuition of 'countervailing incentives' as discussed in Armstrong (2016) would have applied; since a high value consumer also has an equally high value at the other firm, no firm has comparative advantage over any consumer. This lack of pricing advantage leads to marginal cost pricing in equilibrium. But since in my model add-ons are differentiated and there are consumers who value the add-on at firm 1 but not at firm 2, firm 1 has an incentive to raise the add-on price as to exploit this advantage. An intuitive way to think about how this affects equilibrium pricing is dividing the profit function into profit arising from 'total price' and core price. Considering the case of only informed consumers so that $\sigma = 1$, and defining the total product price as $p_T = p_1 + p_U$, (3.4) is rewritten as

$$p_T \left[\frac{1}{4} + \frac{1}{4t}(p_T^* - p_T + 2\rho(w - p_U^*))\right] + p_1 \left[\frac{1}{4} + \frac{1}{4t}(p^* - p_1 - 2\rho(w - p_U^*))\right]. \quad (3.10)$$

Notably, (3.10) is additively separable in p_T and p_1. Firm 1 can therefore effectively consider two separate markets, one for the 'total product' and one for the 'core only product'. This partition reveals the reason for positive add-on markups. Viewing the demand expressions, it is as if firm 1 in the total product market was competing in a Hotelling market where the consumer preference distribution was skewed towards it. Consumers in the 'total product market' for firm 1 generally prefer firm 1 since they would purchase the add-on at firm 1 but may not want to purchase the add-on at firm 2. In contrast, in the core only market, it is as if firm 1 was competing in a Hotelling market where the preference distribution was skewed against it. The difference in demand means that the trade-off between marginal and infra-marginal consumers is different in the two markets and the firms set a higher price in the total product market. Since the add-on price is the difference between the total and core only price, it follows that add-on prices must be positive.

The reason that the two markets can be additively separated is that regardless of chosen prices, each consumer belongs to either the total or core only market. This depends on the assumption of inelastic add-on demands; with elastic add-on demands a price change in the add-on price would make some consumers 'switch markets' and the markets would not be separable in terms of pricing.

3.3.3 Welfare Analysis

Next, consider the issue of social welfare. I define consumer welfare as aggregate consumer utility, that is, the aggregate valuations of purchased products with prices paid and transport subtracted. Social welfare will accordingly be aggregate valuations of purchased products minus transport costs. Production is costless so can be ignored when considering welfare.

With this definition of social welfare, it is straightforward to rank symmetric equilibria in terms of welfare. Supposing first that all consumers were uninformed, under the assumption of a covered core market and symmetric equilibrium all consumers would purchase the core product from the firm they are closest to on the Hotelling line. This is not socially efficient since some consumers in the center of the line ought to choose their furthest away firm in case they prefer the add-on product at that firm. It is socially optimal for such consumers to select the furthest away firm if the created value in terms of add-on surplus is larger than the extra incurred transport cost.

Suppose a consumer is located closer to firm 1 on the Hotelling line, but is an LH type, so prefer the add-on at firm 2. Define then δ as the distance from the center of the Hotelling line at which this consumer must be located to be indifferent between the two firms. This point is found by equating the utilities of choosing each firm in the case where the consumer is closer to firm 1 but has a higher add-on value at firm 2. Equating the two utilities, with the utility for firm 1 on the lefthand side, yields

$$v - p^* - t(\frac{1}{2} - \delta) = v - p^* + w - p_U^* - t(\frac{1}{2} + \delta),$$

where prices of both firms are the equilibrium prices p^* and p_U^*. The distance of this consumer from firm 1 must be half of the Hotelling line plus the additional distance δ. The expression can be rearranged to find the market equilibrium value of δ:

$$\delta = \frac{w - p_U^*}{2t}. \tag{3.11}$$

This value represents the extent of transport cost the consumer is willing to pay to be able to purchase the add-on product at the other firm, under equilibrium prices. However, in terms of social welfare, consumers ought to select their furthest away firm if the total value created is higher than the transport costs incurred. Doing a similar calculation as above, but in terms of social welfare, gives

$$v + w - t(\frac{1}{2} - \delta^*) = v - t(\frac{1}{2} + \delta^*),$$

which can be rearranged to find the socially optimal value of δ, denoted by δ^*:

$$\delta^* = \frac{w}{2t}. \tag{3.12}$$

Comparing (3.11) and (3.12), it can be seen that provided $p_U^* > 0$, there is an inefficient amount of consumers selecting their furthest away firm in the equilibrium. Given the solutions in Proposition 1, this is the case for market equilibrium whenever $\rho > 0$.[13] This is illustrated in Figure 2, which displays the equilibrium amount of selection compared to the socially efficient amount. Furthermore, since if $\sigma > (2t)/(2t+w)$ the add-on price is decreasing in σ, it follows that social surplus is increasing with σ for this range.

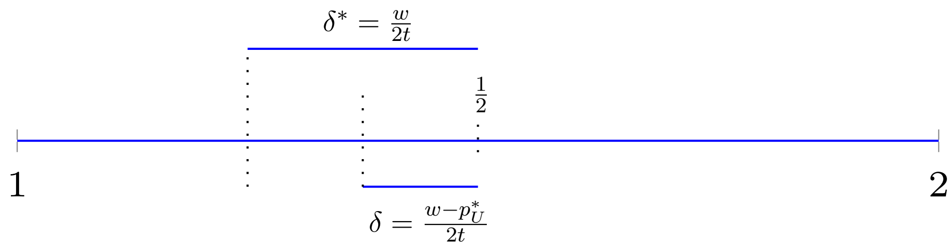

Figure 3.2: Market matching efficiency versus social optimum

Proposition 2 states how social welfare, consumer welfare and firm profits vary with the level of informed consumers. It also gives insight into how welfare effects are distributed between informed and uninformed consumers.

Proposition 2.

(i) If $\sigma < 2t/(2t + w)$, social welfare, firm profits and consumer welfare does not vary with σ.

(ii) If $\sigma > 2t/(2t + w)$, social welfare is increasing in σ, while firm profits are maximised at some $\sigma^* \in (2t/(2t + w), 1]$ and consumer welfare is increasing in σ.

(iii) The ex ante expected utility of uninformed consumers is invariant in σ, while the ex ante expected utility of informed consumers increases in σ when $\sigma > 2t/(2t+w)$.

Proof. See section 3.A.2 of the appendix. □

The proposition states the welfare comparative statics with respect to the level of informed consumers. If there are too few informed consumers there will be monopoly level add-on prices in which case welfare is invariant to changes in the number of informed consumers. This is because in this case the informed consumers are as well

13 Note that $\sigma > 0$ and $\rho = 0$ also gives $\delta < \delta^*$, but in this case there does not exist any consumer that would prefer the add-on product at the furthest away firm since add-on products are then homogenous.

off as the uninformed ones, as all consumers are purchasing the core product from their closest firm regardless of add-on preferences. The high add-on prices mean that consumers receive no surplus from the add-on, so any consumer is indifferent between buying a high-value add-on at monopoly price and not buying the add-on at all. Firms are also in this case indifferent to changes in the number of informed consumers. This is because when $p_U^* = w$, informed and uninformed consumers are in expectation equally profitable to the firms. Note that add-ons are in this case also profit-irrelevant to the firms. The expected profit from the add-on products is fully competed away in core product competition, only leaving the standard Hotelling profit.

When there are sufficient numbers of informed consumers such that add-on prices fall, aggregate social welfare and consumer welfare increases with σ. Part (iii) of Proposition 2 provides insights into the source of this increase. First, the uninformed consumers are indifferent to changes in the level of informed consumers, even though the prices are changing. This result stems from the fact that the add-on and core prices adjust at a relative rate that keeps the total expected price load on uninformed consumers unchanged. In contrast, informed consumers are better off when add-on prices fall. This is because they can now capitalise on their superior information. With a positive add-on surplus to be had, some informed consumers close to the center of the Hotelling line will prefer to switch away from their closest firm due to their add-on preferences. This means that ex ante, informed consumers are more likely to purchase the add-on than uninformed consumers. They therefore care comparatively more about add-on prices than uninformed consumers do and hence benefit when add-on prices fall.

The proof of the proposition shows that equilibrium profits can be written as

$$\pi^* = \frac{1}{2}t + \sigma\rho\delta p_U^*. \tag{3.13}$$

This is intuitive, as profits above the standard Hotelling profit arise only when some consumers select their furthest away firm. This can only happen if the consumer is informed, which is true with probability σ. It must also be the case that the consumer is an HL or LH type, so would be willing to select the furthest away firm, which is true with probability ρ. The consumer must also be placed in the region of the Hotelling line where a consumer would potentially want to select the furthest away firm and this is true with probability δ. Hence the form of the last term in (3.13), which can be read as the probability that a consumer selects firm 1 despite being located closer to firm 2, times the add-on price paid.

To see why the add-on profit is not competed away in the core product market competition, consider the following argument. Suppose $\sigma = 1$ and that both firms are pricing add-ons at p_U^*. Define then $\pi_U(p_1, p^*)$ to be the expected add-on profit per consumer who purchases at firm 1, so the add-on price times the probability that a consumer purchases the add-on at firm 1, conditional on buying the core product at firm 1. Then, the profit of firm 1 can be written as the following, where $D_1(p_1, p^*)$ is shorthand for the demand for firm 1's core product as in (3.4),

$$\pi_1 = [p_1 + \pi_U(p_1, p^*)]\, D_1(p_1, p^*).$$

In this framework, the first order condition with respect to firm 1's core product price when imposing $p_1 = p^*$ becomes

$$p^* = \left(p_1 + \frac{\partial \pi_U(p_1, p^*)}{\partial p_1} \Big|_{p_1 = p^*} \right) t - \pi_U(p^*, p^*).$$

If add-ons were homogenous, the derivative of add-on profitability with respect to core price is zero and the standard result of profit-irrelevant add-ons apply. However, when add-ons are differentiated the expected profitability per consumer is increasing in the core price at firm 1. To see why, note the difference in composition of the firm's marginal and inframarginal consumers. Since some of the inframarginal consumers are HL types that have selected firm 1 due to their add-on preference, the inframarginal consumers are more likely to purchase the add-on at firm 1 than a randomly allocated consumer. In contrast, marginal consumers choosing firm 1 due to a core price cut do not reveal anything about their add-on preferences and are hence no more likely to purchase the add-on than a randomly allocated consumer. Therefore, lowering the core price makes firm 1's consumer base on average less likely to purchase add-ons. This competition softening effect leads to add-on profits not being competed away in the core product market.

3.4 Endogenous Selection

The previous section outlines implications of having a consumer population partially informed about add-ons and how this relates to the differentiation of the add-on products. However, the section did not attempt to explain how consumers become informed and which type of consumers that are most likely to become informed. In particular, the section assumed no correlation between consumers' core product preferences and whether they become informed or not. This section provides a mechanism by which consumers do become informed, arising from rational consumer information acquisition. It is shown that this endogenously generates correlation between consumers core product preferences and becoming informed. I explicitly model the consumers' information acquisition through a costly consumer research process and analyse what implications this has for market equilibrium. It will be assumed that consumers acquire information at both firms simultaneously rather than sequentially. As such, the information gathering process is related to the fixed sample search and clearinghouse strand of the consumer search literature, see Baye *et al.* (2006) for a review of this literature. A justification for this type of search rather than a sequential process may be the existence of a large cost of the first search but low subsequent search costs, so that consumers would optimally either search at every firm or no firm.

The fact that consumers individually select whether to become informed or not have implications for market equilibrium. Intuitively, the consumers with strongest incentives to acquire add-on information are those having weak preferences over core products, i.e. those consumers in the center of the Hotelling line. Therefore, informed consumers end up having weaker firm preferences than those not informed. This section will investigate what this reasoning implies for equilibrium characteristics. The set-up of information acquisition is also used to shed light on the question of whether there are too many or too few consumers endogenously becoming informed in equilibrium.

This section maintains the assumption that consumers have passive beliefs. An analysis of the corresponding model with wary beliefs is given in the appendix section 3.A.6. As in the previous section, there is now a competition softening effect present that causes upward pressures on core prices and corresponding downward, but not fully off-setting, pressure on add-on prices. This also increases consumers' incentives to acquire information, leading there to be more informed consumers for given parameter values.

3.4.1 The research decision

The mechanism by which consumers become informed is as follows. After firms have chosen their prices, consumers observe their core product preferences and the core product prices chosen by firms. Conditional on this information, consumers decide whether to pay a *research cost* $s > 0$, or not. If paying the cost, the consumers can observe the add-on prices at both firms as well as their add-on preferences. After the research decisions are made, the model proceeds as previously. To emphasise that information is now endogenous, I will denote the fraction of informed consumers by a new variable γ (taking the role of σ in the previous section).

When a consumer evaluates whether to pay the research cost, she will compare the cost with the expected utility increment occurring from the research. Define the *default firm* of a consumer to be the firm that the consumer would choose if not investing in research. Then research only results in a utility increment if the research outcome leads to the consumer choosing to purchase from the non-default firm. If the consumer were still to purchase from the default firm, the utility obtained would have been the same as if research had not been conducted, implying no positive utility increment. Since in any symmetric equilibrium consumers believe that the firms are charging identical add-on prices, research is conjectured to lead to a utility increment only if the consumer found a high add-on value at the non-default firm and a low add-on value at the default firm. Therefore, the only rationale for consumers to conduct research is to discover their add-on preferences, not to find out about prices. The fact that consumers still do observe prices as a result of research is crucial for equilibrium outcomes, even if each consumer individually in equilibrium is indifferent to whether price information is observed or not.[14]

It will turn out to be the case that whenever a researching consumer has a high add-on value at the non-default firm and a low add-on value at the default firm, the consumer will choose to purchase from the non-default firm. To see this, suppose that this was not the case so that a consumer continued to purchase from the default firm even if preferring the add-on product at the non-default firm. There is then no utility increment compared to not conducting research and it follows that there was no reason for the consumer to conduct research in the first place. Hence, any researching

14 In fact, if prices were not revealed through the research process the unique equilibrium would be $p_U^* = w$ and no consumer choosing to conduct research.

consumer would indeed choose the non-default firm if preferring the add-on product at that firm. This happens with probability ρ which therefore is the probability of a utility increment occurring. The magnitude of the utility increment will be the difference in utilities obtained at the two firms, conditional on a high add-on value at the non-default firm and a low add-on value at the default firm. Recalling the utility functions in (3.1), the condition for conducting research can be written as the following, with costs on the lefthand side and expected utility increment on the righthand side:

$$
\begin{aligned}
s &< \rho((U_2|w_2 = w) - (U_1|w_1 = 0)), \\
s &< \rho(v - p^* + w - p_U^* - t(1 - \theta) - (v - p_1 - t\theta)), \\
&= \rho(w - p_U^* - p^* - t + 2t\theta + p_1).
\end{aligned} \tag{3.14}
$$

Rearranging to get a condition on the consumer's core product preference, it can be worked out that the consumer conducts research if

$$
\theta > \frac{1}{2} + \frac{1}{2t}(p^* - p_1 - w + p_U^* + s\rho^{-1}) = \underline{\theta}, \tag{3.15}
$$

where $\underline{\theta}$ is defined as the highest value of θ for which a consumer would purchase from firm 1 without conducting research. Note the importance of consumers having passive beliefs, so that changes in core prices do not affect add-on price expectations and therefore research incentives. A consumer will have firm 1 as default firm if $\theta < 1/2 + (1/(2t))(p^* - p_1)$, similar to (3.2) in the previous section. Similarly, an identical calculation shows that a consumer with firm 2 as default firm will engage in research if $\theta < \bar{\theta}$, where $\bar{\theta}$ is defined as the lowest value of θ for which a consumer would purchase from firm 2 without conducting research. This value is

$$
\bar{\theta} = \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + w - p_U^* - s\rho^{-1}). \tag{3.16}
$$

Using these cut-off values of θ, the total amount of researchers in the consumer population can be derived. Defining the length of the interval of the Hotelling line at which consumers will be conducting research as γ^*, this can be calculated as

$$
\begin{aligned}
\gamma^* &= \bar{\theta} - \underline{\theta}, \\
&= \frac{1}{t}(w - p_U^* - s\rho^{-1}).
\end{aligned} \tag{3.17}
$$

Note that here p_U^* represents consumer expectations over the symmetric add-on price at both firms, an expectation which in any equilibrium must be correct and so will be equal to p_U^*. Figure 3 illustrates how the Hotelling Line is partitioned into segments of consumers. The inner segment of consumers will be conducting research and the outer segments will purchase a core product without first conducting research.

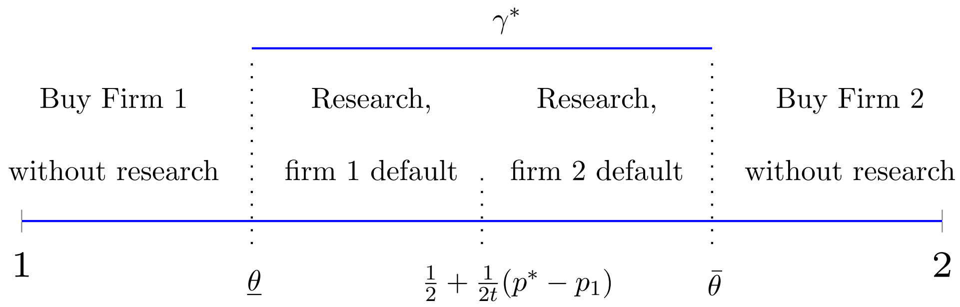

Figure 3.3: Consumer research decision partitioned by Hotelling Line position

Before proceeding, it ought to be remarked that the degenerate equilibrium of having no research at all is always an equilibrium in this model, at any research cost. This is an instance of the Diamond paradox, introduced by Diamond (1971), where search breaks down if consumers believe ex ante that there are no benefits to search, a belief that will be self-enforcing. In my model, if all consumers believe that both firms are charging maximal add-on prices they will find no incentive to conduct research, in which case firms optimally price their add-ons at the maximal price. Noting that this type of equilibrium exists, focus will henceforth be on non-degenerate equilibrium.

3.4.2 Solving for equilibrium

To find an expression for firm profit, consider the various consumer types. First, the amount of uninformed consumers purchasing from firm 1 without research is simply $\underline{\theta}$. These consumers are equally likely to purchase the add-on or not since they have not observed add-on values and the unconditional probability of having a high add-on value is a half. For researchers, the demand for firm 1 arises from the conditional probabilities that the various add-on types are buying from the firm. Defining the events $i \in \{1, 2\}$ as that the consumer buys from firm i, R for the event that the consumer is a researcher and N for the event that the consumer is not a researcher, the profit function for a firm can be written as

$$\pi_1 = (p_1 + p_U)\left[\frac{1}{2}Pr(1,\text{N}) + \gamma^*[(\frac{1}{2} - \rho)Pr(1|\text{R}, HH) + \rho Pr(1|\text{R}, HL)]\right] \qquad (3.18)$$
$$+ \; p_1\left[\frac{1}{2}Pr(1,\text{N}) + \gamma^*[\rho Pr(1|\text{R}, LH) + (\frac{1}{2} - \rho)Pr(1|\text{R}, LL)]\right].$$

To illustrate how the relevant probabilities are found, consider researching consumers that are HH types. Using (3.1), such a consumer will purchase from firm 1 if

$$v - p_1 + w - p_U - t\theta > v - p^* + w - p_U^* - t(1-\theta),$$

which can be rearranged to obtain

$$\theta < \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + p_U^* - p_U). \tag{3.19}$$

Since the consumer is a researcher, the consumer must have $\theta \in [\underline{\theta}, \bar{\theta}]$ and the conditional distribution is uniformly distributed on this interval. Therefore, the conditional probability of (4.6) being true must be

$$\begin{aligned}
Pr(1|R, HH) &= Pr(\theta < \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + p_U^* - p_U)|R), \\
&= Pr\left[\theta \in (\underline{\theta}, \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + p_U^* - p_U))|\theta \in [\underline{\theta}, \bar{\theta}]\right], \\
&= \frac{\frac{1}{2} + \frac{1}{2t}(p^* - p_1 + p_U^* - p_U) - \underline{\theta}}{\bar{\theta} - \underline{\theta}}, \\
&= \frac{\frac{1}{2t}(w - p_U - s\rho^{-1})}{\gamma^*}.
\end{aligned}$$

The last line follows from the definitions of $\bar{\theta}$ and $\underline{\theta}$ from (3.15) and (3.16) and note that only local deviations are considered so that $1/2 + (1/(2t))(p^* - p_1 + p_U^* - p_U) < \bar{\theta}$.

The corresponding probabilities for the other add-on types are of a simple form:

$$\begin{aligned}
Pr(1|R, HL) &= 1, \\
Pr(1|R, LH) &= 0, \\
Pr(1|R, LL) &= \frac{1}{2}.
\end{aligned}$$

These results are justified as follows. Since researching LL types do not buy add-ons, they end up choosing their default firm. Out of the researching population, there is an equal split between those with firm 1 as default firm and those with firm 2 as default firm, so this is also the resulting split for researching LL types. For HL types, their firm preference has been weakly shifted towards firm 1 compared to their prior expectation. Any consumer that had firm 1 as a default firm must therefore purchase from firm 1. Suppose next that the consumer instead had firm 2 as default firm and also chose to purchase from firm 2 despite being an HL type. If this is the case, there is no type realisation that would make this consumer want to switch away from her default firm, meaning that there was no benefit to research in the first place. This is a contradiction, so it must be the case that an HL type purchases from firm 1. Finally, the same logic applies to LH types who must purchase from firm 2 as otherwise, it would not have been rational for them to be researchers in the first place. Note that these arguments rely on that no firm has made a large enough deviation from a conjectured equilibrium where research exists.

Inserting these results into (3.18) and simplifying gives

$$\begin{aligned}\pi_1 = (p_1 + \frac{1}{2}p_U)(\frac{1}{2} + \frac{1}{2t}(p^* - p_1 - w + p_U^* + s\rho^{-1})) \\ + (p_1 + p_U)[\rho\gamma^* + (\frac{1}{2} - \rho)\frac{1}{2t}(w - p_U - s\rho^{-1})] + p_1(\frac{1}{2} - \rho)\frac{1}{2}\gamma^*.\end{aligned} \tag{3.20}$$

The first line refers to profit arising from uninformed consumers and the second line refers to profit arising from researching consumers. It can be seen that there will be a selection effect present similar to the random selection model. Consider the demand for the total (core + add-on) product at firm 1:

$$\rho\gamma^* + (\frac{1}{2} - \rho)\frac{1}{2t}(w - p_U - s\rho^{-1}).$$

In equilibrium this becomes

$$\frac{1}{4}\gamma^*(1 + 2\rho),$$

which can be compared to the equilibrium demand for the core only product from researching consumers, which is given by $(1/4)\gamma^*(1-2\rho)$. The gap of $\rho\gamma^*$ between the demands demonstrates that there are consumers selecting into their preferred add-on firm and demand shifts towards the core plus add-on offering, compared to a random allocation of informed consumers.

3.4.3 Equilibrium Analysis

I now use the profit function (3.20) to derive equilibrium properties. As previously, this is done by letting each firm jointly maximise its profit function by choosing core and add-on prices and from imposing symmetry conditions.

Proposition 3. *Suppose that there exists an interior symmetric equilibrium in pure strategies in which a fraction $\gamma^* \in (0,1)$ conducts research in equilibrium. Then it must be the case that prices are*

$$p^* = t\frac{1-2\rho(2+\gamma^*)}{1-2\rho},$$
$$p_U^* = 4t\rho\frac{1+\gamma^*}{1-2\rho},$$

and the amount of researching consumers is given by

$$\gamma^* = \frac{(1-2\rho)(w-s\rho^{-1})-4\rho t}{t(1+2\rho)}.$$

Proof. If there exists a symmetric equilibrium with $\gamma^* \in (0,1)$, (3.20) is a valid representation of firm profit locally around the equilibrium prices. Taking first order conditions with respect to the core price p_1 and add-on price p_U and imposing equilibrium conditions $p_1 = p^*,\ p_U = p_U^*$ yields the two first order conditions

$$p^* = t - \frac{1}{2}p_U^*,$$
$$p_U^* = t\frac{1+2\rho\gamma^*}{1-2\rho} - p^*.$$

Solving the simultaneous equations yields the solutions given above. Inserting the derived value for p_U^* into the definition of γ^* from (3.17), the given expression for the amount of researching consumers is obtained.[15] □

The proposition states the outcome of the first order conditions of the firms' maximisation problem.[16] The striking feature of the presented results is that add-on prices are in fact increasing in the number of informed consumers. Since core prices move inversely to add-on prices, core prices accordingly decrease with more informed consumers.

The intuition for the reversed price movements in comparison with the random selection model is well illustrated by considering the firms' first order conditions in

15 Appendix section 3.A.3 demonstrates that the second order conditions are satisfied at these prices, given $\rho < 3/7$. To illustrate that there exists parameter values such that simultaneously $p_U^* \leq w$ and $\gamma^* \in (0,1)$, the set of parameter values $t = 1,\ w = 1,\ \rho = 0.05,\ s = 0.005$ gives the following values of the endogenous variables:

$$p^* \approx 0.82,\ p_U^* \approx 0.35,\ \gamma^* \approx 0.55,\ \pi^* \approx 0.51.$$

16 There may be profitable deviations from the candidate equilibrium, for instance when γ^* is very low it would be profitable for firm 1 to deviate to $p_U = w$.

setting the add-on price,

$$0 = t(1-\gamma^*) + t\gamma^*(1+2\rho) - 2(\frac{1}{2} - \rho)(p^* + p_U^*). \tag{3.21}$$

This can be contrasted with the corresponding first order condition from the random selection model

$$0 = t(1-\sigma) + \sigma(t + 2\rho(1-p_U^*)) - \sigma(p^* + p_U^*).$$

The first two terms of both conditions refer to the additional profit from inframarginal consumers when a firm raises the add-on price. This effect is increasing in γ^* or σ as informed consumers are more likely to purchase add-ons, meaning that the mass of inframarginal add-on consumers will be larger. This effect causes upward pressure on add-on prices when there are more informed consumers. The last term in both conditions refers to the profit loss occurring from a price rise as a result of marginal consumers switching firms. Note that γ^* is absent from this effect in the research model, while σ is present in this effect in the random selection model. To see the reason for this note that before simplifying the last term of (3.21) is written

$$2t\gamma^*(\frac{1}{2} - \rho)(p^* + p_U^*)\frac{1}{t\gamma^*}.$$

Viewing the term this way, it is clear that γ^* has two opposing effects on the profit loss on marginal consumers. Firstly, just as in the random selection model, increasing the number of informed consumers increases the number of consumers that react to add-on prices, which explains the lefthand γ^* in the expression. However, with endogenous selection, there is a further effect of increasing the number of informed consumers. As is seen by the second presence of γ^*, increasing the number of informed consumers decreases the sensitivity of add-on buyers to price changes. The reason is that, as is seen in Figure 3.3, the initial researchers are located in the center of the Hotelling line and is therefore very price sensitive. The additionally added researchers when increasing γ^* are less sensitive to prices than these previous researchers. This makes add-on buyers on average less price sensitive which increases incentives for firms to increase add-on prices. It turns out that with a uniform distribution over core preferences this price discrimination effect exactly cancels out with the opposing effect of more consumers reacting to add-on prices.

Because the level of informed consumers cancels out from the effect on marginal consumers, add-on prices must rise with more informed consumers. The striking result of add-on prices increasing with information does depend on the distributional assumptions, e.g. a non-uniform distribution on core preferences may alter this result. The key point is that with endogenous selection there is a price discrimination effect of add-on prices that is not present in the random selection model.

3.4.4 Welfare Analysis

Now consider the welfare properties of the equilibrium with endogenous selection. The main issues of interest are how the welfare properties vary with the research cost s and whether the market provides the socially optimal amount of information acquisition. Welfare is as previously aggregate utility with transport costs and prices paid subtracted and now in addition researching consumers are paying the research cost s. Likewise, social welfare is aggregate utility minus transport costs and minus the research cost paid by researching consumers.

Proposition 4.

(i) Social welfare, firm profits and consumer welfare are all decreasing in s.

(ii) Uninformed consumers are indifferent to the level of s.

(iii) Social welfare had been higher had there been more consumers acquiring information than in the market equilibrium. Consumer welfare had been higher had there been fewer consumers acquiring information.

Proof. See section 3.A.4 of the appendix. □

Part (i) of the proposition states the perhaps not surprising result that social welfare and consumer welfare is lower when research costs are higher. More interesting is the fact that firm profits would also be lower if research costs are increased. This result stems from the fact that with higher research costs, fewer consumers are becoming informed. This means lower levels of firm selection because of add-on values. Just as in the random selection model, abnormal profits above the ones generated by the core differentiation are driven by the fact that some consumers are selecting to purchase from their furthest away firm. This can be seen in the expression for equilibrium profit derived in the proof, which is

$$\pi^* = \frac{1}{2}t + \frac{1}{2}\rho\gamma^* p_U^*.$$

This expression is intuitive. If there was no research, firms would obtain the standard Hotelling profit. With research, the term $\rho\gamma^*$ refers to the probability that a consumer is a researcher and does choose the non-default firm, in which case the firm gains a profit of p_U^*.

The second part of the proposition shows that as in the previous section, uninformed consumers are indifferent to how many informed consumers there are. This is because add-on prices and core prices still adjust at a rate such that the expected payment by uninformed consumers is invariant to changes in γ^*. The last part of the proposition instead shines light on the question of whether there is over or underinvestment in research. The underinvestment in research from a social welfare perspective is because the relevant factors consumers take into account when deciding on research differ from the equivalent social welfare considerations. A consumer deciding whether to conduct research considers the potential benefit from having information, and this benefit is lower when add-on prices are higher. The social welfare benefit from research depends instead on the created value from the potential additional add-on sale and the transport costs incurred. The consumers do not take into account that the prices they pay are profits for firms and is therefore not investing enough into research.

Instead considering consumer welfare, there is overinvestment in research in equilibrium. This stems from the fact that, when consumers decide whether to conduct research, they do not take into account their external effects on the other consumers. Uninformed consumers are indifferent to whether other consumers are informed or not, but informed consumers are worse off when an additional consumer decides to become informed. This is because informed consumers are more sensitive to add-on prices than uninformed consumers, and add-on prices rise when an additional consumer is informed. This last result crucially depends on add-on prices being increasing with the level of informed consumers. As discussed, this feature may not be robust to different assumptions about distribution of preferences.

3.5 Conclusion

This chapter has discussed a duopoly market with add-on products where the add-on products are themselves differentiated independently of the core products offered by firms. Parts of the consumer population are assumed to be uninformed about the add-on products. It has been explained how equilibrium prices are likely to depend on the fraction of consumers that is informed and the level of differentiation in the core and add-on products. The basic model has also been extended to allow for endogenous information acquisition.

Due to a covered market and inelastic add-on demands, the key issue for social welfare is the matching between firms and consumers. If add-on products are differentiated, it is socially optimal for some consumers to select the firm they do not prefer in terms of core products. However, since the equilibrium has above cost add-on prices, there is insufficient add-on surplus offered to attract consumers to the firm they would choose in the social optimum. An implication of this is that the presence of uninformed consumers is particularly harmful when add-on products are differentiated. This since more uninformed consumers lead to higher add-on prices which harms the matching between consumers and firms. The analysis also shows that differentiation in add-ons can be a reason for why firms profit from selling add-ons despite competition in the core product market. The mechanism behind this result is that marginal consumers are less profitable for firms than inframarginal consumers and therefore core price competition is softened.

The analysis of the endogenous research model demonstrates that if consumers were to freely choose whether to acquire add-on information or not, the informed consumer population would consist of those consumers with weakest preferences over firms' core products. This leads to a price discrimination effect where firms have incentives to compete for the informed consumers with low add-on prices, an incentive which diminishes when more consumers become informed. Because of this effect, add-on prices may rise with higher levels of informed consumers. Still, even though add-on prices have risen, welfare increases with informed consumers when the increase in informed consumers arises from reductions in information acquisition costs. Therefore, an intervention which makes information more easily available may have beneficial welfare effects even though the intervention has little discernible downward effect on add-on prices. This is because the improved level of information can lead to improved matching between consumers and firms despite low or even seemingly adverse effects on pricing.

The model generally predicts that lower research costs and more informed consumers are beneficial for welfare. This suggests that interventions making information more easily available will generally be supported given an appropriate cost-benefit assessment. In particular, a novel point is that the presence of uninformed consumers is harmful because of inefficient matching, making a stronger case for intervention than the case of homogenous add-ons. In addition to the basic insight that more transparency is beneficial, the welfare propositions show that there are inter-group effects from improving information that regulators may want to take into account. In particular, uninformed consumers in this model are always unaffected by changes in the availability of information, so any welfare effects only accrues to the informed consumers, or to firms. Finally, one must consider the issue of the relative welfare weights to give to those consumers only buying core products versus those buying both add-on and core products. Whenever the add-on prices rise, core prices fall, meaning that consumers only purchasing the core product benefits, which should be taken into account when making a regulatory intervention.

3.A Appendix

3.A.1 Second order conditions for Proposition 1

The first order partial derivatives of profit function (3.4) with respect to p_1 and p_U are:

$$\begin{aligned}\frac{\partial \pi_1}{\partial p_1} &= (1-\sigma)[\frac{1}{2}+\frac{1}{2t}(p^*-p_1)] - (1-\sigma)\frac{1}{2t}(p_1+\frac{1}{2}p_U)\\ &\quad + \sigma[\frac{1}{4}+\frac{1}{4t}(p^*-p_1+2\rho(1-p_U^*)+p_U^*-p_U)] + \sigma\Big[\frac{1}{4}+\frac{1}{4t}[p^*-p_1-2\rho(w-p_U^*)]\Big]\\ &\quad - \sigma\frac{1}{4t}(p_1+p_U) - \sigma\frac{1}{4t}p_1,\\ \frac{\partial \pi_1}{\partial p_U} &= \frac{1}{2}(1-\sigma)[\frac{1}{2}+\frac{1}{2t}(p^*-p_1)] + \sigma(\frac{1}{4}+\frac{1}{4t}(p^*-p_1+2\rho(1-p_U^*)+p_U^*-p_U)) - \frac{1}{4t}\sigma(p_1+p_U).\end{aligned}$$

Hence, after simplifying, the second order derivatives are

$$\begin{aligned}\frac{\partial^2 \pi_1}{\partial p_1^2} &= -\frac{1}{t},\\ \frac{\partial^2 \pi_1}{\partial p_1 \partial p_U} &= -\frac{1}{4t}(1+\sigma),\\ \frac{\partial^2 \pi_1}{\partial p_U^2} &= -\frac{1}{2t}\sigma.\end{aligned}$$

The Hessian condition then becomes

$$\frac{\sigma}{2t^2} - \frac{(1+\sigma)^2}{16t^2} > 0,$$

which implies

$$\sigma > 3-\sqrt{8} \approx 0.172.$$

This condition will always hold under the conditions $\sigma \geq 2t/(w+2t)$ and $t > w/(1+4\rho)$ that are assumed true. Therefore, the second-order conditions are satisfied.

3.A.2 Proof of Proposition 2

I begin by deriving welfare expressions that will subsequently be used to prove the proposition. Consider first social surplus created from uninformed consumers. Half of these consumers will purchase the add-on product and will in expectation incur a transport cost of $(1/4)t$ since they always choose their closest firm. In addition, since there is a covered market they always purchase the core product to create value v. The social surplus from uninformed consumers is therefore

$$(1-\sigma)(v + \frac{1}{2}w - \frac{1}{4}t).$$

For informed consumers, consider first those consumers with a core product preference strong enough to always purchase from their closest firm. The mass of these consumers is $1 - 2\delta$ in the notation introduced in section 3.3.3. Since these consumers choose their closest firm regardless of add-on preferences, they are as likely to purchase the add-on as uninformed consumers and in expectation create $(1/2)w$ in add-on value. The expected transport cost incurred by these consumers is $(1/4 - (1/2)\delta)t$. Next, consider the consumers that would choose their furthest away firm in case they prefer the add-on at that firm. The mass of these consumers is 2δ. These consumers always buy add-ons unless they are an LL type. Therefore, in expectation they create add-on value $(1/2+\rho)w$. If these consumers purchase from their closest firm, their expected extra transport cost incurred is $(1/2 - (1/2)\delta)t$, which occurs with probability $1-\rho$. If they purchase from their furthest away firm they in expectation incur transport cost $(1/2 + (1/2)\delta)t$.

This yields the following expression for social surplus:

$$\begin{aligned} S = {} & v + (1-\sigma)(\frac{1}{2}w - \frac{1}{4}t) \\ & + \sigma \left[(1-2\delta)(\frac{1}{2}w - (\frac{1}{4} - \frac{1}{2}\delta)t) + 2\delta[(\frac{1}{2}+\rho)w] - [(1-\rho)(\frac{1}{2} - \frac{1}{2}\delta) + \rho(\frac{1}{2} + \frac{1}{2}\delta)]t \right], \end{aligned}$$

which simplifies to

$$v + \frac{1}{2}w - \frac{1}{4}t + 2\sigma\delta\rho[w - t\delta]. \tag{3.22}$$

Notice from (3.4) that equilibrium profits can be written as:

$$\pi^* = (1-\sigma)(\frac{1}{2}p^* + \frac{1}{4}p_U^*) + \sigma[\frac{1}{2}p^* + \frac{1}{4}p_U^* + p_U^*\frac{\rho}{2t}(w - p_U^*)],$$
$$= \frac{1}{2}t + \sigma p_U^* \rho\delta. \qquad (3.23)$$

The second equality is due to the first order conditions of the profit maximisation problem implying $p^* = t - (1/2)p_U^*$ and from the definition of δ.

To obtain the expression for consumer welfare, use the fact that consumer welfare must equal social surplus net of profits. Subtracting (3.23) from (3.22) yields

$$CS = v + \frac{1}{2}w - \frac{5}{4}t + \sigma\delta\rho[w - p_U^*]. \qquad (3.24)$$

The expressions obtained in (3.22), (3.23) and (3.24) are now used to prove the statements made in the proposition.

(*i*) Proposition 1 says that $p_U^* = w$ whenever $\sigma < 2t/(2t+w)$. Equation (3.11) then implies $\delta = 0$ for these values of σ. Inspection of (3.22), (3.23) and (3.24) then reveals that S, π and CS all are independent of σ, and hence the result follows.

(*ii*) First note that when $\sigma > 2t/(2t+w)$, part (*i*) of Proposition 1 implies $p_U^* < w$ which by (3.11) implies $\delta > 0$. Consider the first derivative of (3.22) with respect to σ:

$$\frac{\partial S}{\partial \sigma} = 2\rho\left[\delta(w - t\delta) - \sigma\frac{\partial p_U^*}{\partial \sigma}(w - 2t\delta)\right].$$

Note that $w - 2t\delta = p_U^*$, so this term must be positive. Then, since part (*ii*) of Proposition 1 shows that p_U^* is decreasing in σ, it follows that S is increasing in σ, for this range of σ.

For firm profits, take a first derivative of (3.23) with respect to σ:

$$\frac{\partial \pi}{\partial \sigma} = \rho\left[\delta p_U^* + \frac{\partial p_U^*}{\partial \sigma}\sigma\frac{w - 2p_U^*}{2t}\right],$$

$$\frac{\partial \pi}{\partial \sigma} = \rho\left[\delta p_U^* + \frac{\partial p_U^*}{\partial \sigma}\sigma\frac{w - 2p_U^*}{2t}\right]. \qquad (3.25)$$

This can be shown to be positive whenever p_U^* is larger than $(1/2)w$, which is guaranteed by $\rho > 1/4$. If instead $\rho < 1/4$, the function may be decreasing and there may be an interior profit maximising σ.

For aggregate consumer welfare, the result follows from inspection of (3.24) since $(\partial\delta)(\partial\sigma) > 0$.

(*iii*) Consider the ex ante expected utility for uninformed consumers. Following arguments made when deriving the surplus expressions above, their expected utility is

$$v - p^* + \frac{1}{2}(w - p_U^*) - \frac{1}{4}t.$$

The proof of Proposition 1 shows that the firms' first order conditions imply $p^* = t - (1/2)p_U^*$. Using this, the uninformed consumers' expected utility is

$$v + \frac{1}{2}w - \frac{5}{4}t.$$

The result follows since this expression is independent of σ.

The expected utility of informed consumers is derived similarly as for uninformed consumers, apart from the instance in which they choose their furthest away firm due to add-on preferences. This occurs with probability $2\delta\rho$ and the expected utility increment in this case is $w - p_U^* - \delta t = (w - p_U^*)/2$. Therefore, the expected utility for informed consumers is

$$v - p^* + \frac{1}{2}(w - p_U^*) - \frac{1}{4}t + \delta\rho(w - p_U^*),$$

which is simplified using $p^* = t - (1/2)p_U^*$:

$$v + \frac{1}{2}w - \frac{5}{4}t + \delta\rho\left[w - p_U^*\right].$$

Since δ is increasing and p_U^* decreasing in σ, the expression is increasing in σ. This proves the statement.

3.A.3 Second order conditions for Proposition 3

Using the profit function (3.20), the first derivatives are as follows:

$$\frac{\partial \pi_1}{\partial p_U} = \frac{1}{4} + \frac{1}{4t}(p^* - p_1 - w + p_U^* + s\rho^{-1}) + \rho\gamma^* + (\frac{1}{2} - \rho)\frac{1}{2t}(w - p_U - s\rho^{-1}) - \frac{1}{2t}(\frac{1}{2} - \rho)(p_1 + p_U),$$

$$\frac{\partial \pi_1}{\partial p_1} = \frac{1}{2} + \frac{1}{2t}(p^* - p_1 - w + p_U^* + s\rho^{-1}) + \rho\gamma^*$$
$$+ (\frac{1}{2} - \rho)\frac{1}{2t}(w - p_U - s\rho^{-1}) + \frac{1}{2}(\frac{1}{2} - \rho)\gamma^* - \frac{1}{2t}(p_1 + \frac{1}{2}p_U).$$

Using these first derivatives, the second order derivatives and the cross-derivative can be derived as

$$\frac{\partial^2 \pi_1}{\partial p_1 p_U} = -\frac{1}{4t} - \frac{1}{2t}(\frac{1}{2} - \rho),$$
$$\frac{\partial^2 \pi_1}{\partial p_1^2} = -\frac{1}{t},$$
$$\frac{\partial^2 \pi_1}{\partial p_U^2} = -\frac{1}{t}(\frac{1}{2} - \rho).$$

It can be verified that the Hessian condition becomes $(3 - 7\rho)/(8t^2)$. Therefore, as long as $\rho < 3/7$ the second order conditions are satisfied.

3.A.4 Proof of Proposition 4

(i) An equilibrium social welfare function can be derived similarly to the proof of Proposition 2. For uninformed consumers, the difference to Proposition 2 is that expected transport costs are lower, $(1/4 - (1/4)\gamma^*)t$, since uninformed consumers are located closer to the edges of the Hotelling line. Informed consumers buy an add-on with probability $1/2 + \rho$. With probability ρ the informed consumers choose their furthest away firm and incur expected transport costs $(1/2 + (1/4)\gamma^*)t$. With complementary probability they choose their closest firm and incur expected transport costs $[(1 + \rho)(1/2) - (1 - \rho)(1/4)\gamma^*]t$. Furthermore, the social welfare function takes into account that informed consumers incur a research cost s when choosing to become informed.

The equilibrium social welfare function is then

$$S = v + (1 - \gamma^*)(\frac{1}{2}w - (\frac{1}{4} - \frac{1}{4}\gamma^*)t)$$
$$+ \gamma^*[(\frac{1}{2} + \rho)w - ((1 - \rho)(\frac{1}{2} - \frac{1}{4}\gamma^*) + \rho(\frac{1}{2} + \frac{1}{4}\gamma^*))t - s],$$

which simplifies to

$$v + \frac{1}{2}w - \frac{1}{4}t + \gamma^*(\rho w - \frac{1}{2}\rho\gamma^* t - s). \qquad (3.26)$$

Taking a first derivative with respect to research cost s, and using $(\partial\gamma^*)/(\partial s) = -1/(\rho t)$:

$$-\frac{1}{\rho t}(\rho w - \frac{1}{2}\rho\gamma^* t - s) - \gamma^* + \frac{1}{2}\gamma^*,$$

which simplifies to $\rho t(s - \rho w)$. This expression being positive would imply $\gamma^* < 0$, a contradiction, hence the derivative is negative and welfare is decreasing in research cost s.

For profits, impose equilibrium prices on (3.20) to get

$$\pi^* = \frac{1}{2}p^* + p_U^*(\frac{1}{4} + \frac{1}{2}\rho\gamma^*),$$

which can be simplified using the firms' first order condition $p^* = t - (1/2)p_U^*$:

$$\pi^* = \frac{1}{2}t + p_U^*\frac{1}{2}\rho\gamma^*. \tag{3.27}$$

Since both p_U^* and γ^* are decreasing in s, profits are decreasing in s.

Consumer surplus is found by subtracting firm profits (3.27) from social surplus (3.26):

$$CS = v + \frac{1}{2}w - \frac{5}{4}t + \gamma^*(\rho w - \frac{1}{2}\rho\gamma^* t - s - \rho p_U^*),$$

which simplifies to

$$v + \frac{1}{2}w - \frac{5}{4}t + \frac{1}{2}(\gamma^*)^2\rho,$$

which is clearly decreasing in research costs as s is here only present through γ^*.

(ii) The effect of s on uninformed consumers is through price changes that occur due to changes in s. The expected utility of uninformed consumers is

$$v - p^* + \frac{1}{2}(w - p_U^*) - (\frac{1}{4} - \frac{1}{4}\gamma^*)t.$$

Inserting the firms' first order condition $p^* = t - (1/2)p_U^*$ implies

$$v + \frac{1}{2}w - (\frac{5}{4} - \frac{1}{4}\gamma^*)t.$$

Therefore, the only way research costs affect uninformed consumers is that some previously uninformed consumers become informed. However, those remaining uninformed are indifferent to the research costs.

(iii) Consider a consumer indifferent between researching or not but who chose to be uninformed. This consumer will be located within $(1/4)\gamma^*$ from the center of the Hotelling line. The welfare benefit from informing this consumer arises from the chance that the consumer would find that she preferred the add-on at the furthest away firm. That the consumer would be willing to change firm follows from the indifference between researching or not. The welfare benefit is compared with the additional transport cost $t\gamma^*$ in case of changing firm and the research cost s that would be incurred. It follows that the social welfare impact of informing this consumer is $\rho(w - t\gamma^*) - s$ and using the definition of γ^* this simplifies to ρp_U^*. This is strictly positive, implying that social welfare had been higher had the consumer instead been informed. It follows that there is underprovision of research from a social welfare perspective.

Consider next the impact on consumer welfare from the same exercise. Uninformed consumers are indifferent to an additional consumer becoming informed, as follows from part (*ii*). Since the consumer becoming informed is also indifferent, the only effects must be from informed consumers being affected. The expected utility of a given informed consumer located at θ is

$$v - p^* + (\frac{1}{2} + \rho)(w - p_U^*) - \theta t - s.$$

The firms' first order condition $p^* = t - (1/2)p_U^*$ implies

$$v + (\frac{1}{2} + \rho)w - \rho p_U^* - (\theta + 1)t - s.$$

When informing the additional consumer p_U^* increases according to Proposition 3. Therefore the utility for informed consumers decreases. It follows that consumer welfare would decrease if the indifferent consumer became informed. Conversely, consumer welfare had been higher had the marginal consumer who became informed instead chosen not to become informed.

3.A.5 Wary beliefs in the random selection model

Wary beliefs allow for $p_U^e(p_1) \neq p_U^*$ whenever $p_1 \neq p^*$. When allowing for this, the profit function (3.4) becomes

$$\begin{aligned} \pi_1 = {} & (1-\sigma)(p_1 + \frac{1}{2}p_U)\left[\frac{1}{2} + \frac{1}{2t}[p^* - p_1 + \frac{1}{2}(p_U^* - p_U^e(p_1))]\right] \\ & + \sigma(p_1 + p_U)\left[\frac{1}{4} + \frac{1}{4t}(p^* - p_1 + 2\rho(w - p_U^*) + p_U^* - p_U)\right] \\ & + \sigma p_1 \left[\frac{1}{4} + \frac{1}{4t}[p^* - p_1 - 2\rho(w - p_U^*)]\right]. \end{aligned} \tag{3.28}$$

This is similar to (3.4), but uninformed consumers update their conjecture about add-on prices in response to firm 1's core price. Taking a first derivative of (3.28) with respect to p_1 and imposing symmetric rational beliefs equilibrium requirements $p_1 = p^*, p_U^* = p_U$ and $p_U^e = p_U^*$ gives

$$p^* = \frac{1}{1 + \frac{1}{2}(p_U^e)'(1-\sigma)}t - \frac{1}{2}p_U^*,$$

where $(p_U^e)'$ is the derivative of the consumers' expectation over p_U with respect to core price at firm 1. For the add-on product, each firm maximises (3.28) with respect to their add-on price holding consumer add-on conjectures as fixed. Deriving the first order condition and imposing symmetry as well as rational beliefs gives

$$p_U^* = \frac{t\sigma^{-1} + 2\rho - p^*}{1 + 2\rho}.$$

Solving gives:

$$\begin{aligned} p_U^* &= 2t\frac{(1-\sigma)(1 + \frac{1}{2}(p_U^e)')}{\sigma[1 + \frac{1}{2}(p_U^e)'(1-\sigma)](1+4\rho)} + \frac{4\rho}{1+4\rho}w, \\ p^* &= \frac{\sigma(1+4\rho) - (1-\sigma)(1 + \frac{1}{2}(p_U^e)')}{\sigma[1 + \frac{1}{2}(p_U^e)'(1-\sigma)](1+4\rho)}t - \frac{2\rho}{1+4\rho}w. \end{aligned}$$

Under passive beliefs ($(p_U^e)' = 0$) this is identical to the results from Proposition 1.

The function for expected add-on price is obtained by considering the first order condition for firm 1's add-on price, taking consumer beliefs as given:

$$\begin{aligned} 0 = &\frac{1}{2}(1-\sigma)\left[\frac{1}{2} + \frac{1}{2t}\Big(p^* - p_1 + \frac{1}{2}[p_U^* - p_U^e(p_1)]\Big)\right] \\ &+ \sigma\left[\frac{1}{4} + \frac{1}{4t}(p^* - p_1 + 2\rho(w - p_U^*) + p_U^* - p_U) - \frac{1}{4t}(p_1 + p_U)\right]. \end{aligned}$$

Beliefs being rational implies $p_U^e = p_U$, which means that rational beliefs $p_U^e(p_1)$ become

$$p_U^e(p_1) = \frac{2t + 2p^* - 2(1+\sigma)p_1 + (1+\sigma)p_U^* + 8\sigma t\rho\delta}{(1+3\sigma)}.$$

Most importantly, this reveals that $(p_U^e)' = -(2+2\sigma)/(1+3\sigma) \in [-2, -1]$. Plugging this into the solutions for equilibrium prices gives

$$\begin{aligned} p_U^* &= 4t\frac{(1-\sigma)}{\sigma(3+\sigma)(1+4\rho)} + \frac{4\rho}{1+4\rho}w, \\ p^* &= \frac{(1+3\sigma)(1+4\rho) - 2(1-\sigma)}{\sigma(3+\sigma)(1+4\rho)}t - \frac{2\rho}{1+4\rho}w. \end{aligned}$$

Comparing this with Proposition 1, it can be verified that add-on prices are lower and core prices higher compared to the passive beliefs case. This is because consumers suspect a higher add-on price when a firm cuts its core price, rendering core price cuts less profitable. This raises equilibrium core prices which in turn implies lower add-on prices. It can further be verified that profits are higher with wary beliefs as a result of this competition softening effect in the core price competition, as also pointed out in Armstrong (2015).

It is worth noting that with wary beliefs there exists an interval of σ for which there are multiple equilibria. This is because it is possible that the cut-off in σ at which add-on prices drop from monopoly level is lower with wary beliefs than with passive beliefs. However, if consumers expect $(p_U^e)' = 0$, there is an interval of σ where this will be self-enforcing. Hence, depending on consumer beliefs, multiple equilibria can occur for this region of σ.

3.A.6 Wary beliefs in the endogenous selection model

A useful result regarding information acquisition incentives is the following lemma:

Lemma 1. *Suppose that firm 1 makes a local deviation in its core price. Then, a consumer conducts research if and only if* $\theta \in [\underline{\theta}, \bar{\theta}]$*, where* $\underline{\theta}$ *and* $\bar{\theta}$ *is given by*

$$\underline{\theta} = \frac{1}{2} + \frac{1}{2t}(p^* - p_1 - w + p_U^* + s\rho^{-1}),$$
$$\bar{\theta} = \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + w - p_U^e(p_1) - s\rho^{-1}).$$

The total amount of researchers is then given by

$$\gamma^*(p_1) = \frac{1}{t}(w - \frac{1}{2}(p_U^* + p_U^e(p_1)) - s\rho^{-1}).$$

Proof. Suppose firm 1 has made a local deviation in its core price. Consider first the consumer indifferent between conducting research or buying at firm 2 without research. As previously, the value to research for this consumer is to find out whether she is an HL type. Since this consumer is indifferent between researching and not, the research cost must equal the expected utility increment from research:

$$\begin{aligned} s &< \rho[(U_1 | w_1 = w) - (U_2 | w_1 = 0)], \\ s &< \rho[v + w - p_U^e(p_1) - p^* - t(1-\theta) - (v - p_1 - t\theta)], \\ &= \rho(w - p_U^e(p_1) - p^* - t + 2t\theta + p_1). \end{aligned}$$

This can be rearranged to find

$$\bar{\theta} = \frac{1}{2} + \frac{1}{2t}(p^* - p_1 + w - p_U^e(p_1) - s\rho^{-1}).$$

The same procedure for the consumer indifferent between conducting research or buying at firm 1 without research yields $\underline{\theta} = 1/2 + (1/2t)(p^* - p_1 - w + p_U^* + s\rho^{-1})$. Subtracting $\bar{\theta} - \underline{\theta}$ gives the expression for $\gamma^*(p_1)$. □

The lemma reveals that regarding research incentives, the wary beliefs assumption only directly impacts those consumers with firm 2 as default firm. This is because these consumers are conducting research due to the prospect of purchasing the add-on at firm 1 and therefore take into account their conjecture about firm 1's add-on price when deciding on research. The value to research from consumers with firm 1 as default firm is unaffected by their conjecture about add-on price at firm 1.

The profit function is then given by

$$\begin{aligned}\pi = (p_1 + \frac{1}{2}p_U)\left[\frac{1}{2} + \frac{1}{2t}(p^* - p_1 - w + p_U^* + s\rho^{-1})\right] \\ + (p_1 + p_U)[(\frac{1}{2} - \rho)\frac{1}{2t}(w - p_U - s\rho^{-1}) + \rho\gamma^*(p_1)] + p_1(\frac{1}{2} - \rho)\frac{1}{2t}(w - p_U^* - s\rho^{-1}).\end{aligned}$$

Beliefs are found as previously. Holding consumer beliefs constant, the first-order conditions are

$$\begin{aligned}p^* &= \frac{t}{1 + (p_U^e)'\rho} - \frac{\frac{1}{2} + (p_U^e)'\rho}{1 + (p_U^e)'\rho}p_U^*, \\ p_U^* &= t\frac{1 + 2\rho\gamma^*}{1 - 2\rho} - p^*.\end{aligned}$$

Imposing rational beliefs, $p_U^e(p^*) = p_U^*$, and solving gives

$$p_U^e(p_1) = \frac{t + p^* + p_U^*(1 - 2\rho) - 2s + 2\rho w}{2(1 - \rho)} - p_1.$$

Notably, $(p_U^e)' = -1$ regardless of the level of research costs. Inserting this into the first order conditions given above and solving the set of equations gives

$$\begin{aligned}p^* &= t(1 - 2\rho\gamma^*), \\ p_U^* &= 2t\frac{\rho + 2\rho\gamma^*(1 - \rho)}{1 - 2\rho}.\end{aligned}$$

The equilibrium number of researchers is given by

$$\gamma^* = \frac{(1 - 2\rho)(w - s\rho^{-1}) - 2t\rho}{t(1 + 2\rho - 4\rho^2)}.$$

The qualitative properties of the equilibrium are preserved with wary beliefs, for example add-on prices are still increasing in the number of informed consumers. For given fundamental values, the number of researchers are higher with wary beliefs than with passive beliefs. This is due to the lower add-on prices which increase expected returns to research. Profits are higher with wary beliefs due to the competition softening effect, but firms also benefit from more add-on sales as a result of higher numbers of informed consumers. Another contrast with the passive beliefs model is that with wary beliefs, uninformed consumers are affected when more consumers become informed as firms are no longer trading off the two prices in the same proportion.

4 Competition and Bundling: An Empirical Analysis of European Telecom Markets

4.1 Introduction

In markets where firms are supplying multiple products, it is possible that the firms also offer the products in a joint package to consumers, as a bundle. These bundles may be the only way to acquire the products (pure bundling) or the standalone products may be offered alongside the bundle (mixed bundling). Markets with this feature present challenges to competition policy that are absent in more straightforward uniproduct markets. For instance, under mixed bundling, it may not be clear what the relevant market for the purposes of competition policy is, and it may not be clear how the anti-competitive effects of mergers are affected by the presence of mixed bundling. A market where competition policy and bundling are prominent policy issues is telecom markets. This chapter structurally estimates European telecom demand and utilises the results to investigate implications of mixed bundling for competition policy and market performance. In particular, the chapter investigates welfare effects of counterfactual mergers under mixed bundling and various scenarios with regards to bundling demand.

European telecom markets are increasingly characterised by a mixed bundling market structure. Traditionally telecom products have been viewed as separate product markets, with mainly separate suppliers. This situation has increasingly changed as a result of technological advancement, innovative marketing strategies and cross-market consolidation. It is now common in many countries to offer bundles that combine several types of telecom products. These bundles may consist of any combination of fixed voice, broadband, mobile and TV services (termed single-play, double-play, triple-play or quad-play depending on the number of services bundled). Consumers may be incentivised to purchase these bundles through discounts, additional add-ons or the convenience of only paying one bill. The particular focus of the present chapter is on the bundling of mobile and broadband ('fixed') services. The bundling of these two services has recently been facilitated by increasing convergence between fixed and mobile services and increasing consumer demand for uninterrupted connectivity at all times. This development (termed Fixed-Mobile Convergence or 'FMC' in industry jargon) has been a leading driver behind cross-market consolidation, providing merging firms with the ability to offer both types of services. A further salient characteristic of emerging bundle offers is that the discounts offered by firms can be significant, to incentivise consumers to purchase all their telecom products from the same firm. The uptake of these bundles has in some countries been substantial, such as in Spain and Portugal, while in other countries such as the United Kingdom uptake is lower.

These recent market changes raise non-trivial questions regarding market performance and pose new challenges to competition overview of telecom markets. The fact that the structure of demand is changing has been noted by the European Commission in recent telecom merger investigations, stating that "...the competitive dynamics on both markets are significantly affected by the diminished distinction between fixed and mobile offerings from a demand perspective."[1] The changing market structure could have important implications for competition policy, and misleading inferences could be made if treating telecom products as isolated markets. Previous case precedence may be less informative for future cases in the new market environment, as for example pointed out by Bergqvist and Townsend (2015) regarding the abusive practice of margin squeezes. Competition authorities are yet to take a strong view on the impact of increasing convergence, with the European Commision in several past decisions leaving open the question of whether bundles constitute a separate product market for the purposes of competition policy.[2]

This chapter empirically investigates the impact of bundling on telecom markets, in particular its impact on merger outcomes. This is done by a structural estimation of a model of telecom markets in Europe, which is subsequently used for counterfactual analysis. The model of demand for mobile and broadband services is estimated using two main separate data sets that supply information at different aggregation levels; one at the household level and one at the firm level. The set-up of the model is based on the general framework of the differentiated product demand estimation literature as pioneered by Berry *et al.* (1995), although the model is tailored to the particular

1 Page 39, European Commission (2016a).
2 Pages 20-22, European Commission (2015).

structure of the data sources used. To account for the ability of households to purchase multiple products, a multiple-choice model that is re-cast into a discrete-choice setting is used, following Gentzkow (2007). The model allows for flexible substitution patterns using interaction terms between individual characteristics and product characteristics, as well as through random taste coefficients on the product characteristics. Additional supply-side information and assumptions on pricing behaviour are employed to estimate price sensitivity parameters, in response to the standard endogeneity concerns regarding price coefficients.

The estimated demand model is used together with the assumed supply-side behaviour to simulate counterfactual mergers. I investigate a number of mergers in several markets under different conditions. First, I focus on the role of pricing complementarities between firms' standalone offers for pricing effects of mergers. These complementarities arise as a second firm may lose demand from a price rise at a first firm's standalone product when consumers substitute from standalone products to bundles. I investigate this issue by performing comparable mergers where the merging firms display different patterns in their respective cross-price elasticities. This intends to highlight the importance of complementarity and to what extent a simple market share analysis would be appropriate. Second, given recent trends in demand for bundles I analyse how welfare effects of mergers are affected by the aggregate share of bundles in a market. Therefore, I compute a number of counterfactual mergers first under observed market conditions and then from a counterfactual initial condition where the aggregate market share of bundles is significantly larger.

The estimated results imply that cross-firm complementarities are important in determining pricing effects from a merger. If the merging firms' standalone offerings are complementary, standalone prices may fall as a result of the merger. Bundle prices always rise as a result of a merger and consumer surplus in most estimated cases fall, with some exceptions due to the complementarity effects. The fall in consumer surplus is found to be lower the more complementary the merging firms' standalone offerings. In the presented example, the resulting increase in prices is 6.31 percentage points lower and consumer surplus is \$4.90 per consumer per month larger when the more complementary firms merge, compared to when two less complementary firms merge. It is also found that a market with a larger share of bundles as a result of consumer preferences for bundling have stronger anti-competitive merger effects.

This is because when the merging firms sell more bundles, the firms overall become stronger substitutes. The results indicate that when initial market shares of bundles are increased to match the share in the market with highest bundling uptake, price increases are 1.79 percentage points higher and consumer surplus is on average \$1.61 lower per consumer per month as a result of a merger, in comparison to a merger under observed market conditions.

The remaining part of this section outlines the relevant literature. Section 4.2 describes relevant features of the markets analysed. Section 4.3 describes the data and outlines its features. Section 4.4 outlines the theoretical model, while section 4.5 describes the estimation procedure and discusses identification. The estimation results are presented in section 4.6, while section 4.7 uses these results for counterfactual exercises. Finally, section 4.8 concludes.

Literature

Macieira *et al.* (2013) and Pereira *et al.* (2013) both estimate structural models of demand for telecom bundles. Their preferred demand model is the cross-nested logit, which they find to provide the most reasonable substitution patterns. The former paper analyses the incentives of firms to bundle using consumer level data from the Portuguese telecom industry. Among their findings are that they are unable to find evidence of cost synergies in bundling and that the number of firms present in a local market affects the number of products offered by each firm. Pereira *et al.* (2013) uses consumer level data obtained from Portuguese telecom firms to answer the question of whether telecom bundles constitute a separate product market for the purposes of competition policy. They find that bundles can be seen as a separate market, and find evidence of complementarity between some telecom products. Both these mentioned papers differ from my study by only studying one country-market, Portugal, and both focus solely on triple-play bundles of broadband, fixed phone and television rather than bundling between mobile and broadband services.

Grzybowski and Liang (2015) estimate a model of demand for full quadruple-play bundles (including mobile services). They find evidence of complementarity between mobile data and broadband services, while mobile voice and broadband are estimated to be substitutes. They use data from a single firm in a single town and is therefore unable to conduct analysis of market competition. Grzybowski and Verboven (2016) estimate demand for fixed and mobile voice services in a framework similar to mine and are likewise using the Special Eurobarometer, for the years 2005-2011. Their interest is on fixed voice and mobile voice substitution, but they do not have data on firm-level and so do not comment on competition.[3] There is a larger literature on substitution between fixed and mobile voice services, see for example Grzybowski (2014).

The general set-up of the empirical model builds on the general framework for differentiated product demand estimation as pioneered by Bresnahan (1987), and developed in Berry (1994), Berry *et al.* (1995), Nevo (2000) and Nevo (2001). The model used in this chapter specifically builds upon the multi-product choice framework of Gentzkow (2007), who models consumer choice as a single discrete choice problem by letting consumers choose which collection of products to buy. The principal cost of this modelling convenience is the reliance on a logit error that is independent across product collections and therefore neglects the components that alternatives are made up of. Fan (2013) provides an alternative set-up based on the multiple-discrete choice model by Hendel (1999), but this set-up is not directly transferable to my setting due to my two distinct product groups. In my setting, choosing one firm for broadband rules out all other broadband offers which is not the case in her newspaper setting.

There is a significant empirical literature on merger analysis in differentiated product markets, following Baker and Bresnahan (1985) and Berry and Pakes (1993). Werden and Froeb (1994) provide an analysis of the properties of mergers in a standard logit demand framework. The majority of the literature analyses ex ante counterfactual analysis of potential mergers, exceptions that provide ex post analysis are Peters (2006), Weinberg (2011), Houde (2012) and Bjornerstedt and Verboven (2016). A difficulty in counterfactual merger analysis is the treatment of efficiencies, as empirical evidence by Ashenfelter *et al.* (2015) suggests that synergies can be substantial and outweigh pricing effects from increased concentration. Sometimes an ad hoc 5% rule of cost savings have been used in the literature, for example by Ivaldi and Verboven (2005) in their analysis of the abandoned Volvo-Scania merger. Bonnet and Schain (2017) use input-output data to predict synergy savings and shows significant

3 For two years of the Special Eurobarometer, responses reported whether products were from the incumbent or not, so for this period they do provide an analysis of substitution between incumbent and competitor firms. My model is a variant to this case, but where observed consumer choices are on a more aggregated level and I consider more firms than just a binary incumbent or non-incumbent choice.

heterogeneity, suggesting that such ad hoc rules should be treated with great caution. Genakos *et al.* (2017) analyses mergers in mobile telecom markets and points out the importance of CAPEX investments when considering merger effects. The theoretical literature on mergers in markets with bundling has mainly focused on markets with strictly complementary products. Choi (2008) provides an analysis of complementary markets where mixed bundling is enabled by a merger, motivated by the 2001 GE and Honeywell merger. It is found that the move towards bundling leads to lower prices for bundles but higher prices for mix-and-match systems. Other papers in this strand of the literature include Flores-Fillol and Moner-Colonques (2011) and Mantovani and Vandekerckhove (2016), and a common theme is that if products are complementary, a merger between independent providers leads to lower bundle prices.

Crawford (2008) and Crawford and Yurukoglu (2012) empirically analyses bundling in cable television. The former paper finds evidence of a discriminatory motive for bundling, bundling being a tool for second-degree price discrimination. The latter paper analyses welfare effects of enforcing à la carte pricing and find that rising input costs due to altered bargaining power offsets any benefits from wider consumer choice. The theoretical literature on bundling is more established than its empirical counterpart, after the seminal paper by Adams and Yellen (1976) who showed how a monopolist may strictly benefit from mixed bundling if consumer values are negatively correlated or independent across goods. Most relevant for this chapter is bundling in oligopolies, such as Armstrong and Vickers (2010) and more recently Zhou (2017). The theoretical paper most related to my work is Anderson and Leruth (1993) who present a duopoly model with product differentiation modelled by random logit taste shocks. Economides (1993) shows how mixed bundling can be a Prisoner's Dilemma situation under competition, with firms being better off without bundling but the equilibrium characterised by mixed bundling. Whinston (1990) shows how firms may use bundling to foreclose entry, which has been a prominent concern for antitrust authorities regarding bundling. There is further a number of theoretical papers that consider the effect of the number of firms in bundling markets, including Economides (1989), Kim and Choi (2015) and Zhou (2017).[4]

4 The corresponding working paper version Zhou (2015) analyses mixed bundling.

4.2 European Telecom Markets

In this section, I outline the broad features of the European telecommunication markets as analysed in this paper. I will specifically be analysing two products. The first is mobile services that entail the provision of mobile connectivity through voice and data that can be accessed e.g. through a mobile handset. The second product is internet broadband services that provide internet connectivity to the household. In addition to these services, telecom firms also frequently supply fixed voice services and in some instances TV services, but these market segments are mainly abstracted from in this work. Telecom products are relatively homogenous across firms but may be differentiated in terms of tariff structure, customer services, branding, bundling and add-on content, local variation in network coverage and network quality. Geographically, European telecom markets are distinctly national in scope, due to network infrastructure being national, regulatory regimes being on national level and price offers made on a national basis. In some cases, the coverage of broadband infrastructure may see regional variation, but market definition for the purposes of competition policy has primarily been one of national markets.[5] There is relatively limited cross-country ownership of telecom companies, although multi-market players such as Vodafone and Hutchison are notable exceptions as being present in several European markets.

Traditionally telecom services consisted of fixed voice services supplied by monopolist firms on copper networks. De-regulation followed in the EU in 1998 when entrants started to supply fixed voice services on the incumbent's copper networks and subsequently deployed their own networks. Telecom networks are still regulated with regards to wholesale access prices. Over the last few years, there has been a wave of mergers in telecom markets that has seen the number of competitors of fixed and/or mobile access fall to 3 or 4 in most countries. There have also been relatively few cross-border mergers in recent years, suggesting that gaining scale on a national level is not significantly beneficial to European telecom companies.[6] Many mergers have been between fixed and mobile providers, allowing them to offer both services as part of their bundles. The prominence of bundling has increased significantly in recent years with rising market shares in many markets. Market participants state

5 European Commission (2016b).
6 Telenor's 2014 acquisition of Globul and Germanos in Bulgaria being an exception.

that the most important reason for consumers to choose bundles is price discounts, although convenience and possibility of adding additional products is also a driver of value.[7]

The key players in mobile markets are MNO's (Mobile Network Operators) that run their own mobile network infrastructures. In addition, there are in many markets MVNO's (Mobile Virtual Network Operators) that do not run their own network but purchase wholesale access from network providers. It has been noted in merger investigations that MNO's pose the most significant competitive constraints, with MVNO's competitive power limited in comparison.[8] Both mobile and broadband networks are costly to initiate and maintain, which is one reason for the relatively few number of firms in each market. For mobile networks, the number of firms is also limited by the availability of mobile spectrum frequencies. Furthermore, there are in many markets active network sharing agreements where suppliers share a network in order to save on fixed costs and maintenance costs.

4.3 Data Description

4.3.1 Data

The estimation makes use of data from two main sources, one which provides information at household level and one at firm level. The household level data comes from the 'Special Eurobarometer 438 - E-communications and the Digital Single Market' conducted on behalf of the European Commission and published in 2016. Data collection took place between October 17 and October 26 in 2015 with around 1000 surveys conducted in each country. The survey was conducted in all 28 EU countries and contain responses to questions regarding households' telecom consumption.[9] The key variable regards whether the household is subscribing to a mobile and/or broadband contract and whether the household has purchased these services in a bundle. The data also contains household characteristics such as the age of respondent, community type or occupation as well as some information on telecom habits and usage.

7 See page 21, European Commission (2015).

8 This was, for example, the view expressed by the commission in the Telefonica Deutschland/E-Plus merger enquiry: "Neither MVNO's and Service Providers nor Branded Resellers are at present, or will be after the proposed transaction, able to compete with MNO's in the retail market for mobile telecommunications services in the same way as an MNO". Page 133, European Commission (2014).

9 The survey conducted also had a second topical focus; respondent attitudes to European Common Agriculture Policy.

The sampling methodology is designed so that the survey should be representative on a country level and household respondent is randomly chosen (e.g. by first birthday method). More information about sampling methodology and the full list of survey questions can be found in the Eurobarometer reports on the European Commission website,[10] and the data is publicly available for non-commercial use.

The firm level data used in estimation was compiled by telecom consultancy firm Analysis Mason who compiled publicly available information during September 2016 for 9 countries in the European Union: United Kingdom, Germany, France, Spain, Portugal, Italy, Netherlands, Poland and Romania. Information is reported on prices, market shares and other offer characteristics. The data is collected for all firms in each country that are offering both broadband and mobile services. In some of the countries, there are firms active only in the broadband market or only in the mobile market. I am not able to observe these firms' representative offer characteristics or individual market shares, but can from the data back out their aggregate market share. Firms that supply both mobile and broadband products are denoted 'inside' firms, while a firm that only supplies either mobile or broadband (and hence not a bundle) is denoted an 'outside firm'. In the estimation, I will assume that the product characteristics of the 'outside firm' is the average of the product characteristics of the inside firms.[11] The information that I observe about each firm is on prices, number of active contracts, mobile data allowance, broadband speed, contract duration as well as some further information such as whether there are unlimited mobile calling minutes or the broadband offer includes fixed voice.

The firm level data reports two entries ('advanced' or 'standard') of prices and characteristics for each product type (mobile, broadband or bundle) at each firm and one market share entry. This is an approximation of the true offers facing consumers for a number of reasons; e.g. firms generally offers a number of tariffs at different quality levels and there may at times be significant discounts or special offers present. The data compilation addressed these issues by arguing that the collected 'standard' offers were representative of each firm's mass-market proposition, while the 'advanced' offers are representative of a more premium offer that is still comparable across the

10 ec.europa.eu/commfrontoffice/publicopinion.

11 Furthermore, there are as mentioned in some countries active MVNO's (Mobile Virtual Network Operators) in the mobile market that are renting wholesale network access from mobile network providers. The market share data lists the shares of these companies shares as included in the market share for the mobile provider whose network they access.

firms. Since I observe only one market share observation per firm, I will in estimation only consider the 'standard' offers. This is because these offers are meant to be mass-market propositions and ought to have a larger influence on the market shares than the advanced offers, in particular since market shares are in contract counts rather than e.g. revenue. This type of simplification due to data limitations is widely used, for example Ivaldi and Verboven (2005) only use 'base' models of trucks to represent the full firm offerings.[12] Regarding price data for bundles, there are some data points where I do not observe a bundle price for a pure bundle with only mobile and broadband. Rather, in these cases I observe the price of a bundle that consists of mobile, broadband and fixed voice (and potentially TV services). In these cases, I re-calculate the bundle price to be the sum of the standalone prices adjusted for the aggregate bundle discount that is applied to the larger bundle. This is aimed at realistically capturing the relative price level between standalone and bundle offerings, although comes at the cost of the imputed price not actually being observed.

Regarding mobile offers, all offers in the data set contain some allocation of mobile data. Therefore, in the household data I define that a consumer has chosen mobile if the consumer has a mobile contract that also includes some data services. There are some consumers purchasing mobile services that only include voice services; such consumers will be denoted as having chosen the outside option or broadband only. Furthermore, market shares reported in the firm level data refer to contract (post-paid) shares. Sales of pre-paid or pay-as-you-go services are therefore not included in this data. In the household data, I do not observe whether a given household has pre-paid or post-paid mobile services. I therefore make the assumption that there are no significant deviations between pre-paid and post-paid shares and treat the post-paid market shares as representative of the overall mobile market. In 2013 there was approximately a 50/50 split between pre-paid and post-paid across Europe.[13]

Furthermore, as will be described in section 4.5, my chosen method for estimating the price coefficient requires additional supply-side information. This information comes in the form of accounting data that is suggestive of industry margins. The data used is based on the EBITDA margin as aggregated over countries. I lack such data for Poland, so for this country I use the overall average measure as an approximation. The data was supplied by a consultancy that had collected this data

12 Although concerns about this assumption were raised by Hausman and Leonard (2005).

13 GSMA Intelligence (2013).

from public earnings reports of telecom companies. While accounting data can be a weak approximation of true economic costs and benefits, the implied margins are approximately what one would expect in telecom markets and the estimated price coefficients are not very sensitive to exact margin specification.

4.3.2 Data analysis

This subsection provides key data summary statistics. Table 4.1 describes summary statistics from the firm level data set, in particular the product characteristics. All prices are listed in terms of 2015 PPP adjusted dollars per month. Mobile data is reported in gigabyte per month and broadband speed is in terms of Megabits per second (Mbps). A notable feature of the table is the extent of variation in terms of bundle discounts, both within and across countries.

Next, I focus on the distribution of telecom choices from the household data.[14] The choice alternatives for a household are to not buy either mobile or fixed broadband (0), to buy only mobile (M), to buy only fixed broadband (F), to buy both mobile and fixed but from separate providers (MF), or to buy both from the same provider, that is, buy a bundle (B). Households that are choosing group '0' may be households that are either not accessing telecom services at all, only have mobile voice services or have other means of accessing services than purchasing its own contracts, e.g. through the use of relatives' services. Note that if a consumer is purchasing mobile and broadband from the same firm, I assume that consumer will always purchase the bundle at that firm. This allows me to infer that sales of standalone mobile and broadband refer to sales to consumers that have not purchased a bundle. Without this assumption, market shares would not be identified using my available data.

Table 4.2 displays the choice distribution across household characteristics, averaged across countries.[15] Notably, it is primarily low income and retired households

14 A household is recorded to have broadband if 'internet access at home' is answered as yes to the question D46 'Which of the following goods do you have?'. A household is recorded as having mobile if answering at least one to the question 'Thinking about all household members, including yourself, please indicate for each of the following how many of them are available in your household: Mobile phone subscriptions or pre-paid arrangements giving access to the internet, e.g. for playing or downloading audio or video content or sending and receiving emails'. A household is recorded as having purchased a bundle if answering 'Yes, fixed or mobile internet access' to the question 'By bundle, we mean a package offering a combination of electronic communication services (or a combination with other services such as TV channels) from the same provider at an overall price. Has your household subscribed to two or more of the following services as part of a bundle?'.

15 I had access to further household characteristics, such as respondent occupation and whether the household reported having troubles paying bills. These variables were however not used in the estimation due to statistical insignificance and computational tractability. The same applies to some

Table 4.1: Summary statistics of firm level data

Characteristic	Mean	Within st. dev.	Across st. dev.	Min	Max
Price mobile ($)	29.50	8.56	6.42	2.50	49.54
Price broadband ($)	30.20	5.54	13.33	1.45	61.19
Mobile data (GB)	2.93	2.56	2.51	0.20	40.00
Broadband speed (Mbps)	50.02	17.60	37.55	10.00	200.00
Bundle discount (%)	16.13	6.88	9.68	1.75	29.47
No. of 'inside' firms	3.22	N/A	0.97	2	5

Notes: Overall number of observations (firm-product entries), including outside firms: 102. All reported values exclude outside options but include outside firms. Mobile and broadband prices refer to the relevant standalone offers only. Values for mobile data and broadband speed are computed over both bundles and the relevant standalone offers. Within variation refers to the average of within-country standard deviations. Across variation refers to standard deviation of country averages.

that are choosing the outside option. In particular, 38% of retired households choose the outside options, while this number is only 18% for the full sample. For bundles, households that are large (≥ 3 individuals) or have reported to be of a higher class are most likely to have purchased a bundle. Retired households are less likely to purchase bundles. This may suggest that bundle purchases are not driven by price sensitive consumers looking for discounts, but rather consumers that are either wealthy or with large telecom demand that seek the convenience of only using one firm. Furthermore, a salient feature of the distribution of household choices is substantial variation in bundle market shares across countries. Appendix section 4.A.1 displays the distribution of the household characteristics by country as well as explanations of how each variable is defined.

Table 4.3 displays a regression of product market shares on product characteristics. Price is seen to have a negative influence on market shares but is statistically insignificant from zero. While a negative relationship between price and market share is reasonable, there are endogeneity concerns regarding the relationship between price and market shares, which will subsequently be addressed in estimation. The coefficient on mobile data is negative while the coefficient on broadband speed is positive, although both coefficients are insignificantly different from zero. Finally, when combining the two data sets, there is a concern regarding that the data to some extent contain

product characteristics.

Table 4.2: Distribution of household choices

(%)	0	M	F	MF	B
Overall	17.71	12.64	12.26	42.36	15.04
By Characteristics*					
Hhsizelarge (43.91)	5.74	11.27	9.83	52.13	21.03
Urban (67.30)	15.99	13.29	11.24	43.50	15.97
Lowclass (49.12)	23.78	14.50	13.33	35.94	12.46
Retired (31.52)	38.45	12.52	17.18	25.46	6.39
By Country					
UK	13.27	8.52	15.31	58.60	4.30
FRA	15.27	6.81	14.49	38.81	24.61
GER	17.60	10.73	14.89	47.63	9.15
ITA	20.30	32.12	4.43	31.92	11.23
NET	2.49	1.15	11.59	71.07	13.70
POR	23.60	11.89	19.51	24.36	20.65
SPA	20.30	17.34	5.12	22.76	34.48
POL	20.76	9.48	13.17	44.21	12.38
ROM	27.53	18.39	9.14	33.82	11.11

Notes: Number of household observations: 10083. *Brackets refer to overall share of households that are in this category. Rows sum horizontally to 100.

overlapping information. Since the data sets are from two different sources, these overlapping pieces of information are not necessarily mutually consistent. Therefore, I make certain assumptions to ensure consistency. A full explanation of the issues and the approach taken is explained in Appendix section 4.A.2. One concern is that the firm level data is silent on whether those consumers buying standalone products are only buying that product, or has also purchased another standalone product from a different firm. To make the data sets consistent, I estimate a measure of the number of mobile contracts per household that is allowed to vary by country. The data sets also imply slightly different aggregate level bundle shares. Regarding these bundle shares, I use the aggregate levels from the household data and use the firm data to distribute this share across the firms.

Table 4.3: Market share regressions

Variable / Regression	(1)	(2)	(3)	(4)
Price	-0.08	-0.08	-	-
	(0.06)	(0.06)	()	()
Data	0.02	-	0.02	-
	(0.13)	()	(0.13)	()
BBspeed	-0.01	-	-	-0.01
	(0.02)	()	()	(0.02)
Mobile constant	-3.25	-3.59	-5.11	-4.58
	(1.81)	(1.76)	(1.37)	(1.74)
Broadband constant	-3.98	-3.90	-5.46	-5.37
	(2.72)	(1.74)	(1.41)	(1.34)

Notes: Displayed values are coefficient estimates, with standard errors in brackets. Number of observations (firm-product entries): 102. All regressions include country dummies and a constant. Dependent variable is market share in percentage points.

4.4 Model

The set-up of the model is as follows, where I begin by introducing notation. There is a set of country-markets W indexed by w, a set of firms in each country J_w indexed by j, a set of consumers[16] per country I_w indexed by i and a set of products $K = \{M, F\}$, which are indexed by k. The product M represents the product mobile and F represents (fixed) broadband. I define an *option* r to be a couple (x, y) where x refers to the firm chosen for mobile and y refers to the firm chosen for broadband, where firm 0 represents not purchasing the product at all. The option $(0, 0)$ is therefore the outside option of purchasing neither mobile nor broadband. The options can then be sorted into *groups* G indexed by g which are sets of options. The full set of groups is $\{0, M, F, MF, B\}$ where 0 is the outside option, M is purchasing only mobile, F is purchasing only broadband, MF is purchasing both mobile and broadband but from separate firms and B refers to purchasing a bundle.[17] I assume that if a consumer purchases both products from the same firm then the consumer has chosen a bundle rather than two standalone offers and it then follows that all options are matched into a unique product group.

16 I will use the terms 'consumer' and 'household' interchangeably from here on.

17 So e.g. the option $(0, 2)$ (no mobile, broadband at firm 2) is in group F and the option $(1, 2)$ (mobile at firm 1, broadband at firm 2) is in group MF.

Following Gentzkow (2007), I will write consumer utility as follows, where I drop the country subscript. The utility $\bar{u}_{ijk}$ for a consumer i from consuming a telecom product k purchased from firm j is

$$\bar{u}_{ijk} = \beta_i x_{jk} - \alpha_i p_{jk} + \xi_{jk}, \tag{4.1}$$

where

$$\beta_i = \beta + \beta^o z_i^\beta + \sigma^\beta v_i^\beta,$$
$$\alpha_i = \alpha + \alpha^o z_i^\alpha + \sigma^\alpha v_i^\alpha.$$

Here x_{jk} refers to characteristics of the product and p_{jk} is the price of the product. ξ_{jk} is an unobserved product fixed effect representing unobserved product heterogeneity that is constant across consumers and can be interpreted as a quality index. A key concern in estimation will be potential correlation of this unobserved term with the observed variables, in particular with price, and subsequent sections will deal with this issue. The terms $\{z_i^\beta, z_i^\alpha\}$ represent individual characteristics that help explain preferences over product characteristics and price respectively. The set of $\{v_i^\beta, v_i^\alpha\}$ represents random unobserved taste shocks that are assumed to have a standard normal distribution and to be independently and identically distributed across consumers and products.[18] These are meant to capture unobserved preference variation in preferences over product characteristics (v_i^β) and price (v_i^α).[19] For short, let σ denote the vector of all the coefficients on the random taste shocks. Due to the assumption of standard normal distributions, these estimated coefficients can be interpreted as the standard deviation of the taste shock in the population.

Combining equations, and writing $\delta_{jk} = \beta x_{jk} - \alpha p_{jk} + \xi_{jk}$ to be the mean utility of alternative jk, the utility (4.1) can be written as

$$\begin{aligned}\bar{u}_{ijk} &= \delta_{jk} + \beta^o z_i^\beta x_{jk} + \sigma^\beta v_i^\beta x_{jk} - \alpha^o z_i^\alpha p_{jk} - \sigma^\alpha v_i^\alpha p_{jk}, \\ &= \delta_{jk} + \mu(\beta^o, \alpha^o, \sigma, v_i, z_i, x_{jk}, p_{jk}),\end{aligned}$$

where $\mu(.)$ represents the non-linear part of utility that varies across consumers. Consumers may however decide to purchase more than one standalone product. Therefore, consumer choices are made over options rather than directly on standalone

18 The assumption of normal distribution means that some consumers will with positive probability have preferences that are inconsistent with basic theory, e.g. have a positive coefficient on price. While in a strict sense implausible, the measure of these consumers is in the estimation minor.

19 In the estimation, not all product characteristics will be allowed to have random coefficients due to computational tractability reasons.

products. Consistently with the data, I allow for the possibility that the firm can set a bundle discount when consumers buy a bundle, so that a firm's standalone prices may be higher than the bundled price. Options are indexed by r and let $j(r,k)$ denote the firm choice for product k in option r. Note that this may include the alternative 0 which represents the outside option of not buying the product at all. Next, there is a random taste shock across options for each consumer ϵ_{ir} that is assumed to have an Extreme Value Type 1 distribution, i.i.d. across consumers, yielding the standard logit model choice probabilities. This error term can be interpreted as residual variation in consumer preferences after controlling for observed consumer characteristics. I then write the utility for consumer i of purchasing option r as

$$u_{ir} = \sum_{k \in \{M,F\}} \bar{u}_{ij(r,k)k} + I(r)\Gamma + B(r)\left(\Gamma' + \beta x_r^B + \alpha p_r^B + \xi_r^B\right) + \epsilon_{ir}. \tag{4.2}$$

Where

$$I(r) = 1 \text{ iff } j(r,M) \neq 0 \text{ and } j(r,F) \neq 0,$$
$$B(r) = 1 \text{ iff } j(r,M) = j(r,F) \neq 0.$$

The option utility (4.2) first consists of the sum of the utilities of the underlying products that form the option, while the remaining terms account for portions of utility not present when only considering the standalone offerings. First, there is a product complementarity term Γ that represents the addition to utility that arises when the two products are consumed in conjunction as opposed to consumed in isolation. Next, Γ' is a further addition to utility that arises when both products are purchased at the same firm, i.e. if the option is a bundle. This term may capture the convenience benefit of purchasing both products from the same firm, arising for example through only having to pay one bill or from reduced shopping costs. It may also capture other benefits associated with buying a bundle, such that mobile and broadband may also be bundled with TV or fixed voice services.[20]

The triple $\{x_r^B, p_r^B, \xi_r^B\}$ allows the underlying product characteristics of the bundle to be different from the sum of its underlying components. This because it is observed in the data that some firms' standard bundle offer is different from the sum of its standard standalone offerings. For example, some firms offer bundle discounts or a different (generally higher) level of mobile data allowance for bundles. I also allow

20 As section 4.3.1 mentions, I control for the price impact of these additional services, but the possibility of adding on these services may carry additional utility value.

for bundle-specific unobserved product characteristics, motivated by the fact that many firms also offer the possibility of bundling with fixed voice, cable-TV or other additional benefits to entice buyers. Merging these terms with the sum of the corresponding standalone components results in option characteristics x_r, p_r and ξ_r. Then denote

$$\delta_r = \beta x_r - \alpha p_r + \xi_r + I(r)\Gamma + B(r)\Gamma' \tag{4.3}$$

to be the average utility of an option across consumers. The expression for a consumer's option utility then simplifies to

$$u_{ir} = \delta_r + \mu(\beta^o, \alpha^o, \sigma, v_i, z_i, x_r, p_r).$$

Integrating (4.2) over the extreme value taste shocks and conditioning on (z_i, v_i), the probability that a consumer i chooses firm j for product k is

$$s_{ir} = \frac{exp(\delta_r + \mu(\beta^o, \alpha^o, \sigma, v_i, z_i, x_r, p_r))}{\sum_{r'} exp(\delta_{r'} + \mu(\beta^o, \alpha^o, \sigma, v_i, z_i, x_{r'}, p_{r'}))}. \tag{4.4}$$

4.5 Estimation

4.5.1 Likelihood

The model is estimated using Simulated Maximum Likelihood[21] on data from the 9 countries covered in the firm level data set. A key concern in constructing an estimator is that the estimation needs to take into account that the data sets provide choice information at different aggregation levels and neither at the option-level that the model predicts choices at. The way the estimation procedure approaches this is by constructing a likelihood function that aggregates the predicted model choices to the relevant aggregation level. This estimation procedure allows me to use the two different data sets with different levels of aggregation into one single likelihood function.

The likelihood of observing an outcome of the data, that is an observation of consumer choices over product groups as well as an allocation of firm level market shares, is

$$L = \prod_w \prod_i \prod_r s_{irw}(\delta, \beta^o, \alpha^o, \sigma | z, v, x, p)^{q(i,r)}, \tag{4.5}$$

21 It is known that Simulated Maximum Likelihood is for finite number of simulation draws biased as a result of performing a discrete integration over a non-linear function, see Train (2009). This estimation method is still employed in the literature, for example by Gentzkow (2007) and by Goolsbee and Petrin (2004) who estimate a multinomial probit model using Simulated Maximum Likelihood.

where $q(i, r)$ is an indicator function that is defined to be equal to one if r is in the same group as consumer i's actual choice. The function s_{irw} is defined as in (4.4). Next, as is usual with Maximum Likelihood, I will not maximise the function in (4.5), but rather the log-likelihood which is derived by taking the log of the likelihood function,

$$log(L) = \sum_{w} \sum_{i} \sum_{r} q(i,r) log \left(s_{irw}(\delta, \beta^{o}, \alpha^{o}, \sigma | z, v, x, p) \right).$$

The set of parameters to be estimated include $\{\beta, \alpha, \Gamma, \Gamma', \beta^o, \alpha^o, \sigma, \xi\}$. A key concern for the estimation is the unobserved product heterogeneities in ξ. Given that they enter the model linearly, I am initially not able to separately identify these effects from the other variables that are constant across consumers. Furthermore, there is the standard concern about price endogeneity in that prices may be correlated with these unobserved effects, yielding (upward) biased estimates of the price coefficient. For these reasons, I will estimate the model in three separate stages. First, I will estimate $\{\beta^o, \alpha^o, \sigma, \Gamma\}$ using the likelihood function and as part of this estimation recover the mean utilities δ. Second, I regress the set of mean utilities δ on the linear parameters to estimate $\{\alpha, \beta, \xi, \Gamma'\}$. Finally, I will estimate the price coefficients α_w to be the ones that rationalise the observed prices conditional on the estimated demand system, a supply-side behavioural assumption and aggregated price-cost margin information.

4.5.2 First-stage estimation

Following the procedure used by Berry *et al.* (1995), and widely used elsewhere in the differentiated product demand estimation literature, I maximise the log-likelihood function over the parameters $\{\beta^o, \alpha^o, \sigma, \Gamma\}$. Within this maximisation, the mean utility terms δ will be found for each candidate parameter vector such that the observed market shares are consistent with the ones predicted by the model. It is however not possible to obtain these constants analytically since this would require inverting the market share equations which are highly non-linear. Therefore, the mean utility constants for each firm offering is rather found numerically using a contraction mapping as developed by Berry *et al.* (1995).

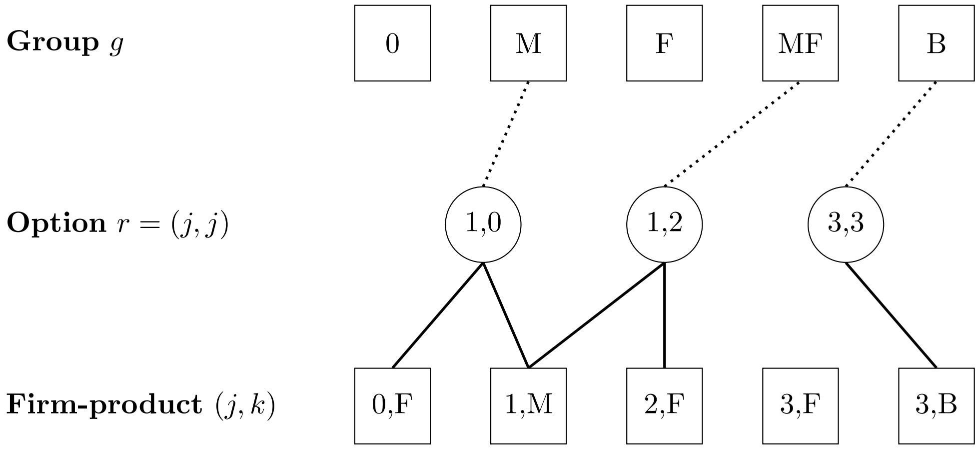

Figure 4.1: Model aggregation levels. Displayed groups are exhaustive, while displayed options and firm-products are examples. Lines represent which firm-product offerings that make up an option and dotted lines represent which group an option belongs to.

The model will yield predictions at the option level, depicted in the middle of Figure 4.1. However, the data reports outcomes only on the group level (household data set) or firm level (firm data set). Therefore, the contraction mapping that will determine the mean utility constants can only identify the mean utility constants at the firm level, not directly identify the constants at the option level. To map mean utilities between firm level and option level the model needs to allow for complementarity effects, so that the mean utility of an option $r = (x, y)$ will be

$$\delta_r = \delta_{xM} + \delta_{yF} + I(r)\Gamma + B(r)\Gamma',$$

where δ_{xM} is the mean utility of mobile at firm x. For this reason, even if Γ is a linear parameter in the sense that it is constant across consumers, I estimate Γ in the first stage estimation. Unlike the other consumer constant parameters, it is identified separately from the unobserved heterogeneities ξ since it only enters utility at the option level while the other linear variables enter at the firm level.[22] Note that if the option is a bundle so that both products are purchased from the same firm, the bundle value Γ' is added. However, since I do have market share information on the bundle at each firm, I am able to directly identify the mean utility of each bundle

22 Note the assumption that Γ is constant across consumers and firms. This assumption is also maintained in Gentzkow (2007), while Grzybowski and Verboven (2016) let this variable be consumer-specific by interacting this coefficient with consumer characteristics.

offering without explicitly adding the Γ'. Therefore, I will not estimate Γ' in the first stage estimation, but rather back this out as part of the second stage estimation.

The predicted market share of an option $\hat{s}_r(.)$ is found by summing over consumers and integrating over the random taste shocks v_i:

$$\hat{s}_r(\beta^o, \alpha^o, \sigma, \delta) = \frac{1}{N} \sum_i \int \frac{exp(\delta_r + \mu(\beta^o, \alpha^o, \sigma, v_i, z_i, x_r, p_r))}{\sum_{r'} exp(\delta_{r'} + \mu(\beta^o, \alpha^o, \sigma, v_i, z_i, x_{r'}, p_{r'}))} \Phi(v_i) dv_i. \quad (4.6)$$

However, the integral in (4.6) does not have a closed-form solution. Rather, the integration is approximated by simulation,[23] so that the computed predicted market share of an option r will be

$$\hat{s}_r(\beta^o, \alpha^o, \sigma, \delta) = \frac{1}{AN} \sum_i \sum_{a=1}^{A} \frac{exp(\delta_r + \mu(\beta^o, \alpha^o, \sigma, v_{ia}, z_i, x_r, p_r))}{\sum_{r'} exp(\delta_{r'} + \mu(\beta^o, \alpha^o, \sigma, v_{ia}, z_i, x_{r'}, p_{r'}))}, \quad (4.7)$$

where v_{ia} is the a'th draw and A is the total number of simulation draws used. Recall that the random draws are from i.i.d. standard normal distributions. The model predictions at the option level $\hat{s}_r(.)$ is aggregated to the firm level $\hat{s}_{jk}(.)$ by

$$\hat{s}_{jk}(\beta^o, \alpha^o, \sigma, \delta) = \sum_{r:j(r,k)=j} \hat{s}_r(\beta^o, \alpha^o, \sigma, \delta),$$

where $\hat{s}_r(\beta^o, \alpha^o, \sigma, \delta)$ is as defined in (4.7).[24] That is, the demand for a product offering is the sum of the demand for the options containing that product. Furthermore, note the abuse of notation such that the subscript k can refer to any of $\{M, F, B\}$, that is, the subscript k can refer to a bundle as the contraction mapping will be run on a firm level including firms' bundles.

The contraction mapping used to find the computed firm level mean utilities $\hat{\delta}_{jk}$ is then

$$\delta_{jkt} = \delta_{jk(t-1)} + log(s_{jk}) - log(\hat{s}_{jk}(\beta^o, \alpha^o, \sigma, \delta_{jk(t-1)})),$$

where $\hat{\delta}_{jk}$ is the fixed point to this mapping. Here s_{jk} refers to the vector of observed market shares and $\hat{s}_{jk}(.)$ to the predicted market shares, all at the firm level. As shown by Berry (1994), there exists a unique set of constants that solves this problem and Berry *et al.* (1995) proved this contraction mapping to be convergent under fairly general conditions.

The set of $\{\beta^o, \alpha^o, \sigma, \Gamma\}$ that maximises the likelihood is the parameter estimates.

23 See Train (2009) for a reference on simulation methods in discrete choice models.

24 Note that these market shares at the firm level sum to more than one due to the presence of two-stop shop consumers who are double counted at the firm level. Similar to the data transformation performed in section 4.A.2, I normalise these market shares to sum to one, interpreting market shares as fractions of household units of each offering.

4.5.3 Second-stage estimation

I now want to recover the linear terms $\{\alpha, \beta, \xi, \Gamma'\}$. This is done in a second-stage OLS regression with the mean utilities $\hat{\delta}_{jk}$ on the lefthand side and the linear variables on the righthand side, in line with (4.3). Therefore, the estimated mean utility for bundles will include Γ that has already been estimated. This is subtracted from the estimated mean utilities so that the lefthand side variable in the second stage regression is interpreted as mean utilities without complementarity effects. Denote $\tilde{\delta}_{jk}$ to be the set of lefthand side variables used in the second stage regression defined as such.

As is frequently argued in the demand estimation literature, there is a worry about the coefficient on price in this regression being endogenous. In particular, the worry is that if there are unobserved product characteristics that affect demand for a product offering, this may affect prices in the same direction as demand and would bias the estimated price coefficient upwards. For example, if a firm has large advertising expenditure this is likely to increase demand but may also increase costs and therefore prices. Other factors causing bias can be brand value or incumbency advantage, leading large firms to set high prices. There are methods to deal with this issue, for example through instrumenting price using appropriate instruments such as cost shifters. I however lack any reliable instruments for the price variable. Therefore, the price coefficient will be estimated in a separate third stage and the second-stage OLS regression will treat price as an unobservable and let the effect from price enter through the constants and the error term. The second stage regression is therefore the following, where product characteristics x include constants for mobile and broadband:

$$\tilde{\delta}_{jk} = \beta x_{jk} + \Gamma' B(r) + (\xi_{jk} - \alpha p_{jk}).$$

4.5.4 Third-stage estimation

Concerned with the problem of price endogeneity, I choose to identify the price coefficient using supply-side assumptions, similar to Smith (2004) and Gentzkow (2007). I assume that firms are choosing prices to satisfy a static Nash equilibrium in prices. The inputs needed to compute the equilibrium include the demand estimation from the first two stages of the estimation as well as a set of marginal costs (three for each firm). However, the firms' marginal costs are unobserved, while actual price choices are observed. Therefore, I will follow common practice in the empirical industrial organisation literature and back out implied marginal costs that follow from imposing Nash Equilibrium conditions on the observed prices and conditional on the estimated demand system.

Under this assumption, each firm will choose prices to maximise its profit function[25]:

$$\pi_j(p) = (p_{jM} - c_{jM})s_{jM}(p) + (p_{jF} - c_{jF})s_{jF}(p) + (p_{jB} - c_{jB})s_{jB}(p),$$

where s_{jk}, c_{jk} refers to demand and marginal cost respectively for product k at firm j and p refers to the vector of all prices in the market. Under the assumption that all firms are setting optimal prices given competitor prices, the first order condition for each firm j when choosing the price of product k is

$$0 = s_{jk} + (p_{jM} - c_{jM})\frac{\partial s_{jM}(p)}{\partial p_{jk}} + (p_{jF} - c_{jF})\frac{\partial s_{jF}(p)}{\partial p_{jk}} + (p_{jB} - c_{jB})\frac{\partial s_{jB}(p)}{\partial p_{jk}}.$$

All first order conditions together define a system $0 = f(\alpha, c)$, where $f(.)$ represents a vector of the righthand side equations as above, which is a function of the price coefficient and the marginal costs. The other relevant factors for the righthand side of the equations are either already estimated or are observed data. In the system of first order conditions, there will be the same number of marginal costs as first order condition equations as each firm has both three prices and three marginal costs. The price coefficients a_w are also unknown, meaning that the system has more unknowns than independent equations and is therefore unidentified. To identify α_w there is therefore additional information needed, which is the accounting data obtained on margins mentioned in section 4.3.1. I employ this information to estimate a separate price coefficient α_w for each country, which means that the system of first order conditions will be just-identified.

For each country w, write the vector of marginal costs c_w as a function of α_w:

$$c_w(\alpha_w) = c_w : f(c_w|\alpha_w, x, p, z, v) = 0.$$

This vector c_w is uniquely defined since the system is just-identified. Furthermore, note that $c_w(\alpha_w)$ is monotonic in α_w. Then, using the margins data μ_w, the estimated

25 One could equivalently write down the profit function with letting the firm choose two standalone prices and one bundle discount.

$\hat{\alpha}_w$ is identified as

$$\hat{\alpha}_w = \alpha_w : \frac{1}{3|J_w|} \sum_{j=1}^{|J_w|-1} \sum_{k \in \{M,F,B\}} \frac{p_{jk} - c_{jk}(\alpha_w)}{p_{jk}} = \mu_w,$$

which uniquely defines $\hat{\alpha}_w$ since $c_w(\alpha_w)$ is monotonic in α_w. Note that the marginal costs are freely estimated without restrictions on the relationship between marginal costs of bundles and the corresponding standalone components. An alternative would be to impose restrictions on this relationship. This however would imply that the model no longer perfectly rationalises prices under observed market conditions, which causes issues with constructing counterfactual exercises. I therefore opt to freely estimate the marginal costs without imposing restrictions.

4.5.5 Identification

I here provide an informal discussion on identification. First, all coefficients apart from the price coefficient are imposed to be identical across countries. A key identifying assumption is therefore stability in preferences across countries. There are also necessary normalisations; the parameters are only identified due to the scaling of utility made by setting the utility of the outside option equal to zero for all countries. All utility values are therefore measured as relative to the outside options, whose utility is set to zero. Furthermore, the scale of the coefficients is normalised such that the variance of the logit error term is equal to one.

For the price coefficient, the identification argument relies on firms playing a static Nash equilibrium in prices and the price coefficent is found to be the one that rationalises the observed prices given the estimated demand model. An advantage of this method is that it is guaranteed to yield a downward sloping demand curve, as theory strongly suggests. The main drawback of the employed method is the reliance on the Nash pricing assumption. This assumption neglects, for example, potential dynamic considerations such as the building of market share or predatory pricing. Caution should therefore be applied to results, as Peters (2006) shows that supply-side misspecification can be a key source causing pricing residuals. However, given the constraints of the data, the static Nash assumption seems a reasonable approximation that has the significant advantage of yielding arguably reasonable model predictions.

The remaining linear variables are estimated through the second stage regression, which is a simple OLS estimation with the estimated mean utilities of each firm-product offering as dependent variable and treating the unobserved heterogeneities ξ_{jk} as a residual. These variables are identified due to variation across firms in terms of the product characteristics and how they vary with the estimated consumer utilities for each firm-product offering. Identification therefore relies on variation observed in the firm level data and on the assumption that the product characteristics are uncorrelated with the unobserved product heterogeneities. This last assumption is in no way innocuous, as I argued the case was with price, and any correlation would bias the estimated coefficient. In theory, one could imagine identifying these coefficients similarly to the price coefficient, for example as described in Crawford (2012). However, given that these estimates will not be crucial for my counterfactuals since they only enter through the mean utilities, I opt for using the second stage regression results.

The identification of the non-linear variables relies on variation observed in the individual level data set and how this varies with the product characteristics in the firm data. This identification is comparatively weak, due to the observed consumer choice being on an aggregated product group level. Identification comes from within countries and from how consumers with different characteristics are choosing product groups which have varying underlying product characteristics. But identification also comes from cross-country variation since the underlying distributions of product characteristics for the product vary by country. The identification of the complementarity parameter Γ estimated in the first stage estimation relies on observing how the market shares of standalone products compare to market shares of multi-stop options and bundles.

Finally, I discuss identification of the random coefficients. The random coefficients allow for flexible substitution patterns by allowing for unobserved consumer characteristics and allowing these to interact with product characteristics. Identification of these coefficients is therefore similar to the coefficients on observed consumer characteristics and arises from cross-country variation in the choice sets. For more discussion of identification of random coefficients in differentiated product market models, see Berry and Haile (2016).

4.6 Results

This section reports the estimates from the preferred model specification. The preferred specification was chosen by considering statistical significance and computational tractability and was estimated using 30 draws from the random coefficient distribution. Table 4.4 reports the estimated coefficients from all three stages of the estimation. The standard errors for the non-linear parameters are the MLE standard errors computed using the finite-difference Hessian as produced by the computational package Matlab. For the linear non-price parameters, the standard OLS standard errors are reported. Since the price coefficients are just-identified, it is not possible to report standard errors for these coefficients.

Considering the non-linear coefficients from the first-stage estimation as displayed in Table 4.4, it can be seen that the price interactions are of the expected signs. Households that have reported to be of a lower social class are significantly more price sensitive than those not reporting to be from a lower social class. Similarly, households where the respondent is retired are significantly more price sensitive than other households. Households that are in urban areas are estimated to have stronger tastes for both broadband speed and mobile data. This seems reasonable given better mobile reception in urban areas and possible complementarity between broadband and mobile coverage. Households with a large number of occupants have stronger tastes for both data and broadband speed. As was suggested in the initial data analysis, retired households have low preferences for bundles while larger households have stronger preferences for bundles. The complementarity parameter Γ is positive and large in magnitude compared to the other coefficients. The random coefficients are significant, suggesting the presence of unexplained preference heterogeneity among consumers. In particular, the random coefficient on bundle preference is large and substantially impacts the implied preference distribution over bundles, which is displayed in appendix section 4.A.3.

Considering the second-stage estimates reported in Table 4.4, the preferred specification treats price as an unobservable left to the error term. This specification is reported in the first column of the second-stage estimates, the second column instead presents regression results where price is included in the regression. If price is included in the second-stage regression, the price coefficient is positive and strongly significant, confirming positive correlation between ξ_{jk} and p_{jk}. Looking at the estimated linear coefficients from the preferred specification in the first column, the coefficients on both mobile data and broadband speed are insignificantly different from zero. However, to properly interpret the coefficients one must consider the impact of the interactions with household characteristics. Section 4.A.3 of the appendix displays the distributions over the population of the coefficient estimates. The coefficient on mobile data is mainly distributed around zero with some consumers having negative tastes for data. The distribution of tastes for broadband speed is mainly positive, with a few households having negative tastes. While the existence of consumers having negative taste for data or broadband speed is implausible, I proceed without further addressing this issue as the level of these coefficients does not impact the counterfactuals performed in section 4.7. Next, Γ' is estimated to be positive, suggesting a shopping cost or utility value of bundling being present. As is displayed in appendix section 4.A.3, there are consumers with implied negative tastes for bundles but they are a small minority.

Table 4.4 also reports the estimated price coefficients, which are estimated separately by country to fit the varying margin data. Recall that prices are measured in PPP adjusted dollars so there is no ex ante expectation of different price coefficients due to aggregate country price levels. All reported values for α_w are positive, which implies downward sloping demands as α enters negatively in the utility function. As with the non-price linear variables, the estimates are only properly interpreted in conjunction with the non-linear estimates. The distribution of the price coefficients across the full consumer sample is displayed in section 4.A.3. There is significant variation in the estimated price coefficients and no consumer has a positive taste for prices, in accordance with basic theory and intuition.

To more clearly interpret the magnitude of these price coefficient estimates, Table 4.5 displays estimated elasticities implied by the model estimates in the Netherlands as an illustrative example. The table provides a set of own-price elasticities as well as cross-price elasticities with two other firms. First, observe that own-price elasticities are varying between -1.64 and -7.16 which is roughly in line with what Pereira *et al.* (2013) find for triple-play markets in Portugal. Cross-price elasticities are very varying, mainly as a result of heterogeneous market shares for the product offerings. However, the most interesting point to be made concern the cross-price elasticities between different firms' mobile and broadband offerings. Most cross-price elasticities

Table 4.4: Estimation results

Variable		Variable		
First-stage		**Second-stage**		
Price*Lowclass	-11.57	Data	-12.98	-7.47
	(0.17)		(33.94)	(32.34)
Price*Retired	-27.02	Bbspeed	0.15	-1.22
	(0.15)		(2.94)	(2.81)
Data*Urban	5.88	Γ'	530.28	773.36
	(0.62)		(379.25)	(364.83)
Data*Hhsize	22.35	Mobile constant	-2884.75	-3768.33
	(1.69)		(239.29)	(290.13)
Bbspeed*Hhsize	4.32	Broadband constant	-2196.85	-2984.85
	(0.05)		(222.46)	(265.46)
Bbspeed*Urban	2.76	Price	-	30.49
	(0.04)		()	(6.19)
Bbspeed*Retired	-1.02	**Third-stage**		
	(0.02)			
		α_{UK}	136.68	
Bundle*Hhsize	45.24			
	(1.19)	α_{FRA}	93.52	
Bundle*Retired	-294.88	α_{GER}	30.24	
	(1.36)			
		α_{ITA}	40.99	
$\sigma_{\Gamma'}$	-101.1			
	(1.80)	α_{NET}	49.96	
σ_{price}	-0.703	α_{POR}	20.28	
	(0.11)			
		α_{SPA}	35.07	
Γ	3564.81			
	(1.01)	α_{POL}	100.93	
		α_{ROM}	65.05	

Notes: Values reported are coefficient estimates with standard errors reported in brackets. All values scaled up three orders of magnitude for expositional clarity.

are positive, due to consumers substituting to other product offerings when a firm raises a price. However, demand for mobile at firm 2 is seen to decrease when firm 1 increases its broadband price. Similarly, demand for broadband at both firm 2 and 3 decreases when firm 1 increases its mobile price. This result does not carry over to the effect on mobile demand at firm 3 when firm 1 increases the broadband price, which instead is increasing.

These results stem from the pricing complementarity created when consumers mix-and-match two firm's standalone options. Consider a consumer that purchases mobile at firm 1 and broadband at firm 2 and that firm 1 increases the standalone mobile price, which induces the consumer to choose a different option. This new option may still include broadband at firm 2, but it may also be to a bundle or to an option with a different broadband provider.[26] Therefore, a price rise at firm 1 can have a negative effect on broadband demand at firm 2. This complementarity effect in prices is weighed against the substitution effect where consumers substitute away from options containing mobile at firm 1 to options that contain broadband at firm 2. In the displayed example it turns out that in one of the cases the complementarity effect is outweighing the substitution effect, but in the other three cases the substitution effect prevails. The implications of this reasoning will be considered in section 4.7.

Table 4.5: Price elasticities of market shares in the Netherlands

	s_{1M}	s_{1F}	s_{1B}	s_{2M}	s_{2F}	s_{2B}	s_{3M}	s_{3F}	s_{3B}
p_{1M}	-3.02	1.16	1.12	1.10	**-1.25**	1.12	1.17	**-1.63**	1.14
p_{1F}	1.49	-2.69	1.43	**-2.04**	1.51	1.43	**-1.18**	1.48	1.46
p_{1B}	0.85	0.85	-6.59	0.87	0.85	0.87	0.85	0.85	0.87
p_{2M}	0.06	**-0.09**	0.06	-5.13	0.06	0.06	0.06	**-0.01**	0.06
p_{2F}	**-1.01**	0.95	0.90	0.88	-1.64	0.90	**0.21**	0.96	0.93
p_{2B}	0.12	0.12	0.13	0.13	0.12	-7.16	0.12	0.12	0.13
p_{3M}	0.59	**-0.46**	0.56	0.55	**0.13**	0.56	-2.44	0.59	0.58
p_{3F}	**-0.08**	0.06	0.06	**-0.01**	0.06	0.06	0.06	-3.28	0.06
p_{3B}	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	-6.28

Notes: Elasticities of demand s_{jk} with respect to price p_{jk} (firm j and product k). Elasticities discussed in text highlighted in bold.

26 A consumer may choose to switch to another standalone broadband provider following a price increase of another firm's mobile offering either because it better complements the consumer's new mobile choice, or because the change of mobile provider leads the consumer to reconsider its broadband choice.

4.7 Counterfactual Simulations

4.7.1 Methodology

With the coefficient estimates in hand, it is possible to use the model to conduct counterfactual simulations. In particular, I will be conducting merger simulations. These merger simulations will be a partial analysis revealing the pricing effects occurring from reduction in competition and ought not to be interpreted as a full evaluation of any conducted merger. This because there are other potential impacts of mergers kept outside of this analysis, e.g. impact on investment incentives as analysed by Federico *et al.* (2017) in general and Genakos *et al.* (2017) for telecoms in particular.[27] There are also potentially synergy and efficiency effects of mergers that are outside of the analysis.[28] As such, any statements about merger effects should be seen as a partial analysis, similar to the suggested interpretation of merger analysis suggested by Baker (1997) and Baker and Rubinfeld (1999).[29]

Regarding methodology for conducting counterfactuals, I will simulate the impact of a merger by assuming that post-merger both merging parties still provide the same separate product offerings but maximise joint rather than separate profits. That is, using the same notation as in section 4.5.3, the merged entity's first-order condition with respect to product k at firm j will be

$$\begin{aligned} 0 = s_{jk} &+ (p_{jM} - c_{jM})\frac{\partial s_{jM}}{\partial p_{jk}} + (p_{jF} - c_{jF})\frac{\partial s_{jF}}{\partial p_{jk}} + (p_{jB} - c_{jB})\frac{\partial s_{jB}}{\partial p_{jk}} \\ &+ (p_{j'M} - c_{j'M})\frac{\partial s_{j'M}}{\partial p_{jk}} + (p_{j'F} - c_{j'F})\frac{\partial s_{j'F}}{\partial p_{jk}} + (p_{j'B} - c_{j'B})\frac{\partial s_{j'B}}{\partial p_{jk}}, \end{aligned}$$

where j and j' are the merging firms. The set of marginal costs here are those predicted by the model, as described in section 4.5.3. This therefore assumes that there are no cost synergies realised by the merger. The set-up also disregards reduction in consumer choice if the merging entity were to offer a single brand, but captures a situ-

27 The direction of this effect is ambiguous, as there may be a trade-off between increased appropriation and reduced investment competition. Federico *et al.* (2017) find that mergers always decrease investment, although the generality of this result has been questioned by both antitrust practitioners and academics, e.g. by Denicolo and Polo (2017).

28 Efficiency savings are often difficult to predict ex ante, although Bonnet and Schain (2017) proposes a model predicting efficiency savings from mergers using input-output data. The European Commission has also generally not looked very favourably on efficiency claims in telecom mergers, see European Commission (2016b). The efficiency savings are often deemed not to satisfy merger specificity and could have been obtained through other means, such as network sharing agreements that are in place in many countries.

29 Baker and Rubinfeld (1999) neatly suggests that merger simulations can be interpreted as "transforming the demand elasticities into a more informative metric".

ation in which the merged entity retains two separate brands. This way of modelling mergers is well established, following Baker and Bresnahan (1985) who also points out that this method is equivalent to modelling a fully collusive agreement between the two participating entities. A further modelling question regards whether post-merger, the mix-and-match options containing products from both merging firms are to be considered a bundle or not. Even if it may be plausible that some of the bundle benefits are present even if two separate brands are maintained, I choose not to add Γ' to the utility of the combined offerings of the merged entities.

Consumer surplus will be defined as the willingness to pay by consumers, less the price paid. Consumer surplus is calculated from the standard property of the logit model which says that (initially shown by Williams (1977) and Small and Rosen (1981)) the total consumer surplus will be

$$CS = \sum_i \int_s \frac{1}{\alpha_{is}} log(\sum_r exp(u_{irs})) f_s ds,$$

where I as in the model estimation will simulate the integral using the simulation draws from the random coefficient distribution.

The profit of a product is naturally the demand of the product times its margin, where margins are based on the estimated marginal costs. A firm's profit is then the sum of the profits of its products. I ignore any fixed costs and welfare measures should therefore be interpreted accordingly.[30] Total welfare straightforwardly consists of the sum of firm profits and consumer surplus.

4.7.2 Mergers - pricing complementarity

This section simulates a set of counterfactual mergers that illustrate the impact of pricing complementarities between firms' standalone products. The mergers conducted are in Portugal, which is relatively representative in terms of the number of firms (3 inside firms) and uptake of bundles. Market shares are relatively evenly spread in the market across both firms and products, and as such is a suitable example of mergers in mixed bundling markets. The firm-product price elasticities are displayed in Table 4.6. As can be seen, there is a similar pattern to Table 4.5 in that many cross-product elasticities are negative, in fact in this case all of them are negative.

30 Neglecting fixed cost does not impact the welfare assessment of the merger, under the assumption that mergers do not affect fixed costs. Even if the merger affects fixed cost, this ought not to affect consumer surplus as only marginal cost changes are passed through to consumers.

However, looking at the magnitudes it appears that they are generally more strongly negative between firm 1 and 2 than between firm 2 and 3. For example, the elasticity of mobile demand at firm 1 with respect to broadband price at firm 2 is -0.88, while the elasticity of mobile demand at firm 3 with respect to the same price is -0.20. The magnitudes of the cross-product elasticities between firm 1 and 3 generally lie in between the other two cases. According to the reasoning presented in the end of section 4.6, a merger between firm 1 and 2 ought then be less anti-competitive than a merger between firm 2 and 3.

Table 4.6: Price elasticities of market shares in Portugal

	s_{1M}	s_{1F}	s_{1B}	s_{2M}	s_{2F}	s_{2B}	s_{3M}	s_{3F}	s_{3B}
p_{1M}	-1.76	0.46	0.47	0.47	**-1.00**	0.48	0.48	**-0.93**	0.48
p_{1F}	0.22	-0.91	0.23	**-0.59**	0.22	0.23	**-0.39**	0.22	0.21
p_{1B}	0.34	0.35	-2.60	0.34	0.34	0.35	0.34	0.34	0.34
p_{2M}	0.23	**-0.62**	0.24	-1.99	0.23	0.24	0.23	**-0.19**	0.23
p_{2F}	**-0.88**	0.41	0.42	0.42	-1.77	0.42	**-0.20**	0.42	0.42
p_{2B}	0.28	0.29	0.29	0.29	0.28	-2.72	0.28	0.28	0.29
p_{3M}	0.28	**-0.48**	0.28	0.28	**-0.13**	0.28	-2.30	0.28	0.29
p_{3F}	**-0.25**	0.13	0.13	**-0.11**	0.13	0.13	0.13	-2.14	0.13
p_{3B}	0.09	0.09	0.09	0.09	0.09	0.09	0.10	0.09	-3.78

Notes: Elasticities of demand s_{jk} with respect to price p_{jk} (firm j and product k). Elasticities discussed in text highlighted in bold.

Table 4.7 displays the impact on consumer surplus and profits from the three possible mergers in the market. As one may expect from a market concentration, prices increases in all cases as a result of the merger. It is also clear that it is the 2 & 3 merger that has the most adverse effects on consumer welfare, with demand-weighted prices increasing by nearly 7%. Saliently, both the 1 & 2 merger and the 1 & 3 merger increases consumer surplus, while the 2 & 3 merger decreases consumer surplus.[31] The outcomes are in line with the predictions that the merger between the less complementary firms leads to a more anti-competitive outcome. The 2 & 3 merger leads to a 6.31 percentage points larger increase in market prices and results

31 Consumer surplus can increase despite rising average prices, as long as some prices are falling, since consumers can substitute their consumption choices in response to the new market prices.

in \$4.90 less surplus per consumer per month, compared to the 1 & 2 merger.[32]

Table 4.8 displays more detailed merger results by describing the movements in prices and demand as a result of the merger. Looking at the 1 & 2 merger, the model predicts that both the standalone mobile and broadband prices at firm 2 would fall as a result of the merger. For firm 1, the broadband price would rise while the mobile price would fall. In contrast, the bundle prices for the merging firms always increase in each of the mergers. This is intuitive since bundles are strictly substitutable to each other and a merger reduces competitive pressure on bundles. Comparing the 1 & 2 merger to the 2 & 3 merger, in the former merger all prices increase as a result of the merger, suggesting its stronger anti-competitive effects. In particular, the non-merging firm increases prices for the 2 & 3 merger, while the non-merging firm reduces its prices in response to a 1 & 2 merger.

Table 4.7: Merger welfare results Portugal

Merger	CS_{pre}	CS_{post}	ΔCS	$\Delta CS(\%)$	$\Delta\pi(\%)$	$\Delta p(\%)$
1 & 2 merger	59.82	61.82	2.00	3.34	-8.24	0.50
1 & 3 merger	59.82	60.43	0.61	1.02	-1.26	0.82
2 & 3 merger	59.82	56.92	-2.90	-4.85	4.55	6.81

Notes: Consumer surplus reported in terms of \$ per consumer per month and profit changes refer to aggregate market profits. Price changes are based on demand-weighted market prices, where weights are based on initial market shares.

32 It can be observed that some mergers lead to reduced market profits. If the profit change for the merging firms is negative, it ought not to endogenously occur unless there are reasons for merging outside the scope of this analysis, such as efficiencies or facilitating collusive outcomes.

Table 4.8: Merger results Portugal

	$1M$	$1F$	$1B$	$2M$	$2F$	$2B$	$3M$	$3F$	$3B$
1 & 2 merger									
$\Delta p\%$	-13.36	4.43	24.53	-22.03	-12.42	41.75	-3.35	-4.56	-2.18
$\Delta s\%$	25.69	-2.67	-46.19	22.48	19.80	-61.44	-2.63	12.08	-8.92
1 & 3 merger									
$\Delta p\%$	2.04	-13.78	9.30	-1.55	-0.68	-0.53	1.54	-11.06	30.86
$\Delta s\%$	2.26	6.46	-20.05	11.49	-5.49	-3.18	1.93	11.81	-55.64
2 & 3 merger									
$\Delta p\%$	2.33	3.54	1.53	8.46	4.52	11.40	10.30	19.30	20.40
$\Delta s\%$	0.93	-2.58	7.17	-5.30	-1.69	-11.86	-9.05	-20.79	-34.03

4.7.3 Mergers - Bundling prominence

Next, I am interested in answering the question of how the prominence of bundles impacts the pricing effects from mergers. I simulate mergers under two different scenarios. First, a merger is simulated under actual market conditions. Second, I compute a new market equilibrium where I increase consumers' preference for bundling Γ' such that the new equilibrium aggregate share of bundles is equal to the aggregate share in the market with the highest observed share of bundles (Spain).[33] The mergers selected for this exercise are examples of mergers between the two largest firms in a market,[34] for select markets. The examples are the discussed cases of Portugal and the Netherlands as well as France and the UK, who differ in the number of inside firms in the market (4 and 5 respectively).

33 The addition to bundling utility to achieve the required equilibrium ranges from 0.64 for Portugal to 2.62 for the UK. For comparison, the mean of Γ' is 0.46.

34 Largest in the sense of aggregate market share across products.

Table 4.9: Bundle demand mergers

Merger	CS_{pre}	CS_{post}	ΔCS	$\Delta CS(\%)$	$\Delta\pi(\%)$	$\Delta p(\%)$
NET - Actual	74.70	64.13	-10.57	-14.14	19.80	13.67
NET - High bundle	78.15	65.93	-12.22	-15.64	19.71	18.42
POR - Actual	59.82	61.82	2.00	3.34	-6.25	0.50
POR - High bundle	65.55	63.46	-2.09	-3.19	-2.71	1.84
UK - Actual	13.41	12.90	-0.51	-3.80	6.87	2.59
UK - High bundle	14.97	14.26	-0.71	-4.74	8.64	3.28
FRA - Actual	21.77	21.17	-0.60	-2.76	3.56	2.16
FRA - High bundle	23.36	22.25	-1.11	-4.75	4.59	2.51

Notes: Consumer surplus reported in terms of $ per consumer per month and profit changes refer to aggregate market profits. Price changes are based on demand-weighted market prices, where weights are based on initial market shares.

Table 4.9 displays the results from these simulated mergers. It is clear from the results that in each case the effect of the merger is more anti-competitive when market conditions are such that bundles are in high demand. In the higher bundling case, the increase in demand-weighted prices is on average 1.79 percentage points higher in the counterfactual case of higher bundling demand. The fall in consumer surplus is larger both in absolute value and in percentage terms under the high bundling demand case for each presented merger. On average for the presented mergers, the merger results in $1.61 less surplus per consumer per month in the higher bundling demand counterfactual than under actual conditions. Next, Table 4.10 displays the price and demand movements resulting from the counterfactual merger between firm 1 and 2 in the Netherlands.[35] The price increases are slightly larger in the higher bundling demand case, in particular for firm 1. The contrast is also clearly seen from that firm 3 would decrease its prices in response to the merger under actual conditions, while in the higher bundling demand case firm 3 would increase its prices as a result of the merger.

35 It can be noted that merger effects appear large in the Netherlands compared to the other markets. This is likely due to a combination of relatively few firms, asymmetric market shares and a small market share of the outside option.

Table 4.10: Bundle demand mergers the Netherlands

	$1M$	$1F$	$1B$	$2M$	$2F$	$2B$	$3M$	$3F$	$3B$
Actual									
$\Delta p\%$	-15.45	36.27	20.36	-10.62	64.84	53.01	-0.62	-2.22	-0.94
$\Delta s\%$	19.40	-13.53	-25.07	14.71	-8.79	-79.52	-22.23	162.86	88.49
High bundle									
$\Delta p\%$	-13.60	38.59	20.25	-6.37	87.02	61.30	7.13	1.68	1.61
$\Delta s\%$	32.42	-6.36	-18.90	17.21	-8.71	-82.91	-20.58	202.56	101.54

4.8 Conclusion

This chapter has presented a model to estimate demand for bundles with data sets providing choice information at different aggregation levels. The model has been estimated using data from European telecom markets, exploiting cross-country variation as well as variation across consumers in survey data to increase model flexibility. Crucially, the consumer level data only provides choices at the product group level rather than at firm level. While the model has been designed with this issue in mind, the lack of information regarding consumers' firm preferences creates challenges for identifying substitution patterns. Finding substitution patterns is also made difficult by the data being cross-sectional, meaning that no substitutions even at the group level are actually observed. Rather, substitution patterns are identified through cross-country variation in choice sets. With these data limitations in mind, the model generates reasonable flexibility, as the cross-price elasticities do display some variation. A potential alternative approach to these issues would be estimating a nested logit or a model with a more general GEV structure, although the random coefficients used approximate such structures, see discussion in Brownstone and Train (1998) and McFadden and Train (2000).

The counterfactual mergers suggest that potential pricing complementarity between firms' standalone offers can be important for predicted pricing effects of mergers. In particular, if consumers tend to consume two different firms' standalone offers in conjunction, a merger between these two firms may lower the prices of their standalone products and increase consumer surplus even in the absence of efficiencies. In the presented example of two comparable mergers that differ in their initial cross-price elasticities, it is found that the resulting price increase is 6.31 percentage points higher for the less complementary of the two mergers. In terms of consumer surplus, this translates to a difference between the merger outcomes of \$4.90 per consumer per month. This heterogeneity in impact suggests that between-firm complementarity is important for assessing pricing effects of mergers. A word of caution is in place regarding the result that some mergers are beneficial for consumer surplus. An implication of the logit errors at the option-level of the model is that without sufficient flexibility, the model is likely to overpredict the amount of substitution from standalone products to standalone products at other firms. In a purely standard logit model without flexibility, substitution patterns are in proportion to market shares of the different options. However, it may be reasonable to believe that a consumer would be more likely to substitute to options that contain the same standalone choice. This means that the model may to some extent overestimate the complementarities in pricing between the firms. However, the comparative merger exercises ought still to be informative of how heterogeneity in the analysed issues can impact outcomes.

In contrast to standalone products, the results indicate that bundle prices always rise after a merger. This is natural, since bundles are strict substitutes to each other and a merger therefore relieves competitive pressure. This intuition leads to the insight that if bundles are more prominent in a market, the anti-competitive effects of mergers ought to be greater. This intuition is demonstrated by the presented simulations, which show that the resulting increase in prices is on average 1.79 percentage points higher in the counterfactual scenario where the aggregate market share of bundles was at the level of Spain, the market with the highest observed bundle uptake in the sample. The effect translates to that the negative impact on consumer surplus is \$1.61 per consumer per month greater in the counterfactual with higher bundling demand.

The results presented in the chapter follows the well established insight in the competition economics literature that purely relying on market shares when assessing anti-competitive pricing effects from mergers may not provide an accurate assessment. In particular, the results illustrate particular issues to pay attention to when assessing mergers in markets with mixed bundling. As such, while the model is estimated with regards to telecom markets, these insights are of importance for any merger occurring in markets which are characterised by mixed bundling. For telecom markets in particular, the results suggest that authorities may want to pay attention to projected developments in bundle uptake, as merger outcomes may depend on the level of market shares of bundles. Since larger prominence of bundles may lead to further competition concerns outside the scope of the present analysis, such as reduced consumer switching, this is an important issue to consider.

4.A Appendix

4.A.1 Household characteristics summary

Table 4.11: Household characteristics variable descriptions

Variable	Description
Hhsizelarge	There are 3 or more residents in the household
Urban	Household is NOT in a ‘rural area or village’
Lowclass	Respondent identifies as ‘working class’ or ‘lower middle class’
Retired	Occupation of the respondent is reported as ‘retired’

Table 4.12: Summary statistics of household characteristics

%	UK	FRA	GER	ITA	NET	POR	SPA	POL	ROM
Hhsizelarge	39.29	37.26	30.54	54.09	35.44	50.90	47.78	51.30	57.62
Urban	67.19	80.64	66.81	86.70	60.44	67.84	59.70	65.07	51.62
Lowclass	88.31	44.75	38.30	34.78	19.25	74.60	65.32	34.93	35.79
Retired	38.84	35.31	35.39	24.24	30.17	26.83	24.83	26.85	36.87

Notes: Number of household observations: 10083. Entries refer to percentage of households in the country that is entered as 1 for the variable.

4.A.2 Data Consistency

The fact that both the individual and the firm level data set contain information related to market shares, albeit at different aggregation levels, raises the issue of need for data consistency. To get sensible estimates, I will assume that both data sets are generated from the same data generating process and in line with this assumption make a number of impositions. The first issue concerns what units market shares ought to be calculated on. The natural way to write down market shares in the firm data would be shares of contracts sold. In the individual data, the natural way to write down market shares is in terms of share of households. To ensure consistent and comparable measures, I use share of households as the unit that market shares are defined in terms of. This makes market shares straightforward to calculate in the individual data set. The market share of product group g will then be the share of households that have chosen that product group. The more involved part is translating market shares from the firm data into comparable market shares in terms of household units.

First, consider the aggregate market share of product group 'O' of buying neither a broadband or mobile contract. This share is observed in the individual level data set. In the firm data, this share is unobserved since the information given is purely in terms of actual sales, and not in terms of potential market. Hence there is no potential for inconsistency across the data sets and I impose that the market share of the outside good will be the one observed in the individual level data set.

Some notation to aid this section: For each product group $g \in G = \{0, M, F, MF, B\}$, say that g_w^I is the absolute count of choices in the individual level data set that belong to group g. For example, if 238 households in France chose mobile only, then $g_{\text{France}}^I = 238$. Similarly define g_w^F to be the absolute observed count of contracts in the firm level data set that belong to group g. For example, if 300 000 standalone mobile contracts were sold in Romania, then $M_{\text{Romania}}^F = 300\ 000$.

Next, consider how to translate the contract sales data into household units. The first step consists of ensuring that the mobile and broadband sales data are comparable in terms of household units. For this, I need the average number of mobile contract purchases per each broadband contract for a general household in each country. This is found by deriving the implied ratio of households that have purchased mobile versus broadband in each data set and equating these measures. For the individual level data, this measure is simply

$$\frac{MF_w^I + M_w^I}{MF_w^I + F_w^I}. \tag{4.8}$$

This measure may not be transferred into household counts before first knowing how many mobile contracts that an average household would have for each broadband contract it purchases. Denote this measure to be a_w. The implied fraction of households that are purchasing standalone mobile versus standalone broadband in each country in the firm data is

$$\frac{M_w^F}{a_w * F_w^F}. \tag{4.9}$$

To ensure consistency, I then impose (4.8) = (4.9), which defines a_w to be

$$a_w = \frac{M_w^F(MF_w^I + F_w^I)}{F_w^F(MF_w^I + M_w^I)}.$$

The estimated a_w is reported in Table 1, where it can be noted that the average is around 2 which seems reasonable, although there are outliers.

Table 4.13: Estimated mobile contract sales per broadband contract sale by country

	UK	France	Germany	Italy	Netherlands	Portugal	Spain	Poland	Romania
a_w	2.30	2.63	2.04	0.77	2.30	3.26	1.96	4.42	2.17

The remaining potential source of data set inconsistency is the aggregate level of market shares for bundles. The implied levels differ slightly across the data sets, but to ensure consistency I opt to use the household data to determine the aggregate level of bundle market shares and use the firm data to distribute this share across the firms. Finally, note that care must be taken in translating contract count data into market shares in terms of household, even after having accounted for a_w as described above. To see this, let $\tilde{M}_w^F = \frac{M_w^F}{a_w}$ to be the adjusted mobile contract count so that now the measure refers to household units and consider using the simple addition $\tilde{M}_w^F + F_w^F + B_w^F$ to represent the total number of households purchasing inside options. This would be misleading and overstate the number of households since some households are two-stop shopping mobile and broadband and so would be double-counted. Therefore, I make use of the observed distribution across product groups MF, M, F in each country to control for double-counting. This allows for correct translation of market shares between data sets.

4.A.3 Taste parameter distributions

Table 4.14: Distribution of coefficients pooled over countries

Coefficient	Mean (standard deviation)	Range
Price	78.00 (42.17)	[18.36,177.23]
Mobile Data	0.79 (11.37)	[-12.98,15.25]
Broadband speed	3.59 (2.70)	[-0.87,7.24]
Γ'	455.67 (178.91)	[-191.08,943.17]

Notes: All values scaled up three orders of magnitude for expositional clarity.

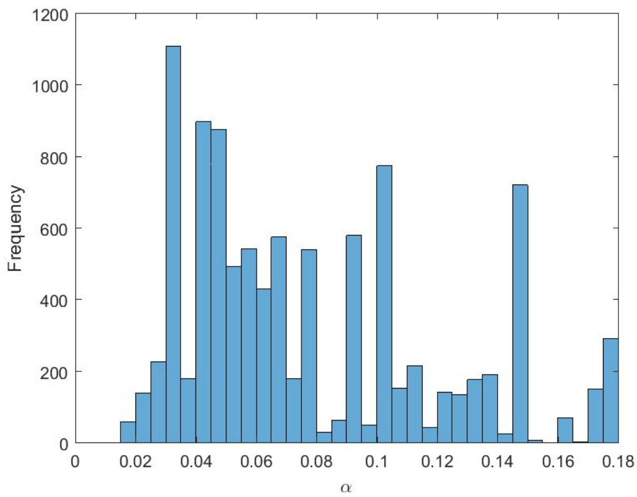

Figure 4.2: Distribution of α

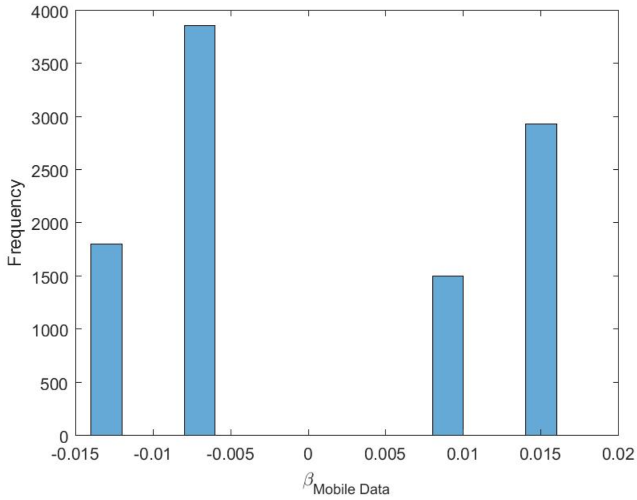

Figure 4.3: Distribution of β_{data}

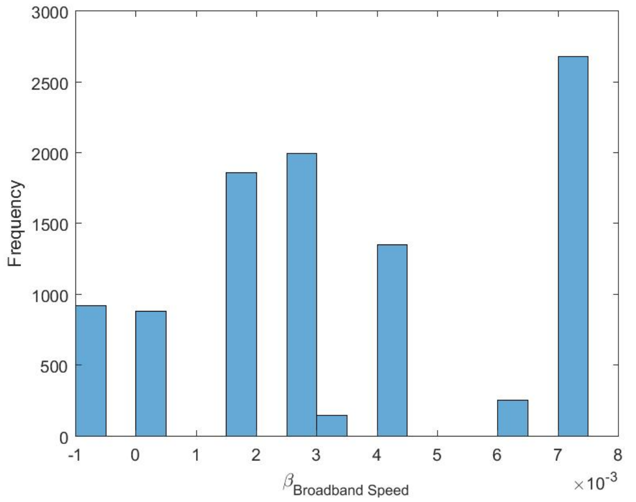

Figure 4.4: Distribution of $\beta_{bbspeed}$

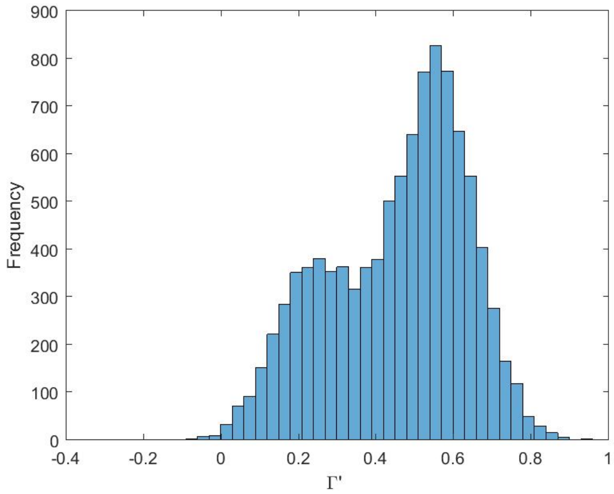

Figure 4.5: Distribution of Γ'

Bibliography

Adams, W. and Yellen, J. L. (1976). Commodity bundling and the burden of monopoly. *The Quarterly Journal of Economics*, **90** (3), 475–498.

Aghion, P. and Hermalin, B. (1990). Legal restrictions on private contracts can enhance efficiency. *Journal of Law, Economics and Organization*, **6** (2), 381–409.

Anderson, S. and Leruth, L. (1993). Why firms may prefer not to price discriminate via mixed bundling. *International Journal of Industrial Organization*, **11** (1), 49–61.

Armstrong, M. (2015). Search and ripoff externalities. *Review of Industrial Organization*, **47** (3), 273–302.

— (2016). Nonlinear pricing. *Annual Review of Economics*, **8** (1), 583–614.

— and Vickers, J. (2001). Competitive price discrimination. *RAND Journal of Economics*, **32** (4), 579–605.

— and — (2010). Competitive non-linear pricing and bundling. *Review of Economic Studies*, **77** (1), 30–60.

— and — (2012). Consumer protection and contingent charges. *Journal of Economic Literature*, **50** (2), 477–493.

Arya, A. and Mittendorf, B. (2004). Using return policies to elicit retailer information. *The RAND Journal of Economics*, **35** (3), 617–630.

Ashenfelter, O. C., Hosken, D. S. and Weinberg, M. C. (2015). Efficiencies brewed: pricing and consolidation in the US beer industry. *The RAND Journal of Economics*, **46** (2), 328–361.

Auster, S. (2013). Asymmetric awareness and moral hazard. *Games and Economic Behavior*, **82**, 503–521.

Baker, J. B. (1997). Contemporary empirical merger analysis. *George Mason Law Review*, **5**, 347–361.

— and Bresnahan, T. F. (1985). The gains from merger or collusion in product-differentiated industries. *The Journal of Industrial Economics*, **33** (4), 427–444.

— and Rubinfeld, D. L. (1999). Empirical methods in antitrust litigation: Review and critique. *American Law and Economics Review*, **1** (1-2), 386–435.

Baye, M. R., Morgan, J., Scholten, P. *et al.* (2006). Information, search, and price dispersion. *Handbook on economics and information systems*, **1**, 323–375.

Bergqvist, C. and Townsend, J. (2015). *Enforcing margin squeeze ex post across converging telecommunications markets*. Konkurrensverket Working Paper Series in Law and Economics, Konkurrensverket.

Berry, S. and Haile, P. (2016). Identification in differentiated products markets. *Annual Review of Economics*, **8** (1), 27–52.

—, Levinsohn, J. and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, **63** (4), 841–890.

— and Pakes, A. (1993). Some applications and limitations of recent advances in empirical industrial organization: Merger analysis. *American Economic Review*, **83** (2), 247–252.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, **25** (2), 242–262.

Bjornerstedt, J. and Verboven, F. (2016). Does merger simulation work? Evidence from the Swedish analgesics market. *American Economic Journal: Applied Economics*, **8** (3), 125–164.

Bonatti, A. (2011). Brand-specific tastes for quality. *International Journal of Industrial Organization*, **29** (5), 562–575.

Bonnet, C. and Schain, J. P. (2017). *An empirical analysis of mergers: Efficiency gains and impact on consumer prices*. TSE Working Paper 244, Toulouse School of Economics.

Bresnahan, T. F. (1987). Competition and collusion in the american automobile industry: The 1955 price war. *Journal of Industrial Economics*, **35** (4), 457–482.

Brownstone, D. and Train, K. (1998). Forecasting new product penetration with flexible substitution patterns. *Journal of Econometrics*, **89** (1-2), 109–129.

Choi, J. P. (2008). Mergers with bundling in complementary markets. *The Journal of Industrial Economics*, **56** (3), 553–577.

Crawford, G. (2008). The discriminatory incentives to bundle in the cable television industry. *Quantitative Marketing and Economics (QME)*, **6** (1), 41–78.

Crawford, G. S. (2012). Endogenous product choice: A progress report. *International Journal of Industrial Organization*, **30** (3), 315–320.

— and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, **102** (2), 643–685.

Dekel, E., Lipman, B. L. and Rustichini, A. (1998). Standard state-space models preclude unawareness. *Econometrica*, **66** (1), 159–174.

Denicolo, V. and Polo, M. (2017). *Duplicative research, mergers and innovation.* CEPR Discussion Paper 12511, CEPR.

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

Economides, N. (1989). Desirability of compatibility in the absence of network externalities. *The American Economic Review*, **79** (5), 1165–1181.

— (1993). *Mixed Bundling in Duopoly.* Stern School of Business Discussion Paper EC-93-29, Stern School of Business.

Ellison, G. (2005). A model of add-on pricing. *The Quarterly Journal of Economics*, **120** (2), 585–637.

European Commission (2014). *Commission decision of 2.7.2014.* Case M.7018 - Telefonica Deutschland / E-Plus.

European Commission (2015). *Commission decision of 19.5.2015.* Case M.7421 - Orange / Jazztel.

EUROPEAN COMMISSION (2016a). *Commission decision of 1.9.2016*. Case M.7758 - Hutchison 3G Italy / WIND / JV.

EUROPEAN COMMISSION (2016b). Recent developments in telecoms mergers. *Competition Merger Brief*, **3/2016**.

FAN, Y. (2013). Ownership consolidation and product characteristics: A study of the US daily newspaper market. *American Economic Review*, **103** (5), 1598–1628.

FEDERICO, G., LANGUS, G. and VALLETTI, T. (2017). A simple model of mergers and innovation. *Economics Letters*, **157** (C), 136–140.

FEINBERG, Y. (2012). *Games with Unawareness*. Research Paper 2122, Stanford University, Graduate School of Business.

FILIZ-OZBAY, E. (2012). Incorporating unawareness into contract theory. *Games and Economic Behavior*, **76** (1), 181–194.

FINANCIAL CONDUCT AUTHORITY (2014). *MS14/1 General insurance add-ons: Final report - confirmed findings of the market study*. Retrieved from Financial Conduct Authority, URL: fca.org.uk/your-fca/documents/market-studies/ms14-01-final-report. Last accessed 17-04-2018.

FLORES-FILLOL, R. and MONER-COLONQUES, R. (2011). Endogenous mergers of complements with mixed bundling. *Review of Industrial Organization*, **39** (3), 231–251.

GABAIX, X. and LAIBSON, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, **121** (2), 505–540.

GENAKOS, C. D., VALLETTI, T. and VERBOVEN, F. (2017). *Evaluating Market Consolidation in Mobile Communications*. CEPR Discussion Paper 12054, CEPR.

GENTZKOW, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, **97** (3), 713–744.

GOOLSBEE, A. and PETRIN, A. (2004). The consumer gains from direct broadcast satellites and the competition with cable TV. *Econometrica*, **72** (2), 351–381.

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

Grzybowski, L. (2014). Fixed-to-mobile substitution in the European Union. *Telecommunications Policy*, **38** (7), 601–612.

— and Liang, J. (2015). Estimating demand for fixed-mobile bundles and switching costs between tariffs. *Information Economics and Policy*, **33** (C), 1–10.

— and Verboven, F. (2016). Substitution between fixed-line and mobile access: the role of complementarities. *Journal of Regulatory Economics*, **49** (2), 113–151.

GSMA Intelligence (2013). *The mobile economy 2013*. Retrieved from GSMA Intelligence, URL: gsma.com/newsroom/wp-content/uploads/2013/12/gsma-mobile-economy-2013.pdf. Last accessed 17-04-2018.

Halpern, J. Y. and Rego, L. C. (2014). Extensive games with possibly unaware players. *Mathematical Social Sciences*, **70**, 42–58.

Hausman, J. and Leonard, G. K. (2005). Using merger simulation models: Testing the underlying assumptions. *International Journal of Industrial Organization*, **23** (9-10), 693–698.

Heidhues, P., Koszegi, B. and Murooka, T. (2017). Inferior products and profitable deception. *The Review of Economic Studies*, **84** (1), 323–356.

Heifetz, A., Meier, M. and Schipper, B. C. (2006). Interactive unawareness. *Journal of Economic Theory*, **130** (1), 78–94.

—, — and — (2013). Unawareness, beliefs, and speculative trade. *Games and Economic Behavior*, **77** (1), 100–121.

Hendel, I. (1999). Estimating multiple-discrete choice models: An application to computerization returns. *The Review of Economic Studies*, **66** (2), 423–446.

Hintikka, J. (1962). *Knowledge and Belief*. Ithaca, N.Y., Cornell University Press.

Houde, J.-F. (2012). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review*, **102** (5), 2147–2182.

IVALDI, M. and VERBOVEN, F. (2005). Quantifying the effects from horizontal mergers in European competition policy. *International Journal of Industrial Organization*, **23** (9-10), 669–691.

JOHNSON, J. P. and MYATT, D. P. (2003). Multiproduct quality competition: Fighting brands and product line pruning. *American Economic Review*, **93** (3), 748–774.

— and — (2006). Multiproduct cournot oligopoly. *RAND Journal of Economics*, **37** (3), 583–601.

KARNI, E. and VIERØ, M.-L. (2017). Awareness of unawareness: A theory of decision making in the face of ignorance. *Journal of Economic Theory*, **168**, 301–328.

KHORANA, A., SERVAES, H. and TUFANO, P. (2009). Mutual fund fees around the world. *Review of Financial Studies*, **22** (3), 1279–1310.

KIM, S.-H. and CHOI, J. P. (2015). Optimal compatibility in systems markets. *Games and Economic Behavior*, **90** (Supplement C), 106–118.

LAL, R. and MATUTES, C. (1994). Retail pricing and advertising strategies. *The Journal of Business*, **67** (3), 345–370.

LI, J. (2009). Information structures with unawareness. *Journal of Economic Theory*, **144** (3), 977–993.

LIU, H. and OZER, O. (2010). Channel incentives in sharing new product demand information and robust contracts. *European Journal of Operational Research*, **207** (3), 1341–1349.

MACIEIRA, J., PEREIRA, P. and VAREDA, J. (2013). *Bundling Incentives in Markets with Product Complementarities: The Case of Triple-Play*. Working Papers 13-15, NET Institute.

MALKIEL, B. G. (2013). Asset management fees and the growth of finance. *Journal of Economic Perspectives*, **27** (2), 97–108.

MANTOVANI, A. and VANDEKERCKHOVE, J. (2016). The strategic interplay between bundling and merging in complementary markets. *Managerial and Decision Economics*, **37** (1), 19–36.

Maskin, E. and Tirole, J. (1990). The principal-agent relationship with an informed principal: The case of private values. *Econometrica*, **58** (2), 379–409.

— and — (1992). The principal-agent relationship with an informed principal, II: Common values. *Econometrica*, **60** (1), 1–42.

McAfee, R. P. and Schwartz, M. (1994). Opportunism in multilateral vertical contracting: Nondiscrimination, exclusivity, and uniformity. *American Economic Review*, **84** (1), 210–230.

McFadden, D. and Train, K. (2000). Mixed MNL models for discrete response. *Journal of Applied Econometrics*, **15** (5), 447–470.

Modica, S. and Rustichini, A. (1999). Unawareness and partitional information structures. *Games and Economic Behavior*, **27** (2), 265–298.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, **18** (2), 301–317.

Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, **31** (3), 395–421.

— (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, **69** (2), 307–342.

Pereira, P., Ribeiro, T. and Vareda, J. (2013). Delineating markets for bundles with consumer level data: The case of triple-play. *International Journal of Industrial Organization*, **31** (6), 760–773.

Peters, C. (2006). Evaluating the performance of merger simulation: Evidence from the U.S. airline industry. *The Journal of Law and Economics*, **49** (2), 627–649.

Rochet, J.-C. and Stole, L. A. (2002). Nonlinear pricing with random participation. *Review of Economic Studies*, **69** (1), 277–311.

Schipper, B. (2013). Awareness-dependent subjective expected utility. *International Journal of Game Theory*, **42** (3), 725–753.

Small, K. A. and Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica*, **49** (1), 105–130.

Smith, H. (2004). Supermarket choice and supermarket competition in market equilibrium. *The Review of Economic Studies*, **71** (1), 235–263.

Stango, V. and Zinman, J. (2014). Limited and varying consumer attention: Evidence from shocks to the salience of bank overdraft fees. *Review of Financial Studies*, **27** (4), 990–1030.

Tirole, J. (2009). Cognition and incomplete contracts. *American Economic Review*, **99** (1), 265–294.

Train, K. E. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press, 2nd edn.

Vasconcelos, L. (2014). Contractual signaling, relationship-specific investment and exclusive agreements. *Games and Economic Behavior*, **87** (C), 19–33.

Verboven, F. (1999). Product line rivalry and market segmentation–with an application to automobile optional engine pricing. *Journal of Industrial Economics*, **47** (4), 399–425.

Weinberg, M. C. (2011). More evidence on the performance of merger simulations. *American Economic Review*, **101** (3), 51–55.

Werden, G. and Froeb, L. (1994). The effects of mergers in differentiated products industries: Logit demand and merger policy. *Journal of Law, Economics and Organization*, **10** (2), 407–426.

Whinston, M. (1990). Tying, foreclosure, and exclusion. *American Economic Review*, **80** (4), 837–59.

Williams, H. C. W. L. (1977). On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and Planning A: Economy and Space*, **9** (3), 285–344.

Wilson, R. (1993). *Nonlinear Pricing*. Oxford University Press.

Zhou, J. (2015). *Competitive Bundling*. MPRA Paper 68358, University Library of Munich.

— (2017). Competitive bundling. *Econometrica*, **85** (1), 145–172.