The London School of Economics and Political Science

Essays on Consumer Learning and Behavioural Economics

Pedro Sequeira Alves

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy, London, July 2016

To Joana.

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of about 24,221 words.

Acknowledgement

I would like to express my deepest gratitude to Gilat Levy for her support and patience. As my supervisor, she provided invaluable insights and guidance throughout the dissertation process. For their generosity with their time and for making me think deeper about my research, I am grateful to Erik Eyster, Francesco Nava and Matthew Levy.

I owe thanks to Jane Dickson and Mark Wilbor for their help throughout the process.

I would also like to thank my collegues and classmates, including Michel Azulai, Alexia Delfino, Florian Blum, Shiyu Bo, Carlo Cabrera, Miguel Espinosa, Dana Kassem, Panos Mavrokonstantis, Marcus Roel, Francesco Sannino, Roberto Sormani, Manuel Staab, Arthur Seibold, Munir Squires, Eddy Tam, Kostas Tokis, Torsten Walter, Lisa Windsteiger, Junichi Yamasaki, Giulia Zane and many others, for their useful comments and their friendship.

I am grateful for the support of Fundacao para a Ciencia e Tecnologia (FCT) for providing the funding necessary to produce this research[1].

The burden of writing this thesis was lessened substantially by the support and humor of my ex-roommates Linda Zhang and Christina Ammon, and the "cool office".

I would like to express my eternal gratitude to my parents for all their patience, unconditional love and encouragement in this PhD journey.

Finally, to my number one fan (Joana), I have no words to express how grateful I am and promise to thank her every day (possibly by taking her to all those exotic trips we have been dreaming about).

[1] Through research grant SFR/BD/73978/2010

Abstract

From its inception, behavioural economics' mission has been to bring deeper psychological insights into economics. Relying mostly on experimental data, this field became notorious for providing evidence of the shortcomings of standard economic models in predicting human behaviour. These findings motivated a first generation of behavioural models, which tried to systematise this departure from standard economics. However, these initial attempts were widely criticised for their methods (these models were argued to lack the tractability, systematic approach and level of generality desired by economic science) and for their lack of relevance for economic phenomena (markets, evolution and arbitrage would drive away behavioural biases). This criticism motivated a second wave of behavioural models, which augmented neo-classical frameworks with psychologically realistic behavioural assumptions. This approach allowed this field to establish a link to previous results of economics and address criticisms about the relevance of behavioural findings in markets.

A further step in the direction of linking behavioural models and standard theory is to introduce learning to behavioural models. While this concept has been largely absent from behavioural economics' analysis of markets for technical reasons, its presence is necessary for two reasons. First, learning is commonly used to dismiss (behaviourally motivated) consumer mistakes, so it is crucial to study whether existing results of this literature will be robust to this variation. Second, in a world which is constantly evolving, learning in itself is an important driver of economic phenomena and, hence, should not be dismissed by this field.

In this thesis, I augment previous behavioural models by studying their existence in environments with consumer learning. By extending static behavioural problems to dynamic environments with learning, I am able to explain puzzles in the areas of technology adoption and contract theory. In chapter 1, I propose that status considerations – a feature of consumers' preferences overlooked by classical theory – can have positive effects in society whenever they are considered in an environment

with active learning (i.e., experimentation). In chapter 2 and 3, I show that when naïve of behavioural consumers (who lack self-awareness about their preferences) can learn, pricing methods in subscription contracts, which were previously unexplained by standard contract theory, can be shown to be the optimal response of firms trying to prevent consumer learning.

Contents

List of Figures

Preface

From its inception, behavioural economics' mission has been to bring deeper psychological insights into economics. Relying mostly on experimental data, this field became notorious for providing evidence of the shortcomings of standard economic models in predicting human behaviour. These findings motivated a first generation of behavioural models, which tried to systematise this departure from standard economics.

However, these initial attempts were widely criticised for their methods (these models were argued to lack the tractability, systematic approach and level of generality desired by economic science) and for their lack of relevance for economic phenomena (markets, evolution and arbitrage would drive away behavioural biases).

This criticism motivated a second wave of behavioural models, which augmented neo-classical frameworks with psychologically realistic behavioural assumptions. This approach allowed this field to establish a link to previous results of economics and address criticisms about the relevance of behavioural findings in markets.

A further step in the direction of linking behavioural models and standard theory is to introduce learning to behavioural models. While this concept has been largely absent from behavioural economics' analysis of markets for technical reasons, its presence is necessary for two reasons. First, learning is commonly used to dismiss (behaviourally motivated) consumer mistakes, so it is crucial to study whether existing results of this literature will be robust to this variation. Second, in a world which is constantly evolving, learning in itself is an important driver of economic phenomena and, hence, should not be dismissed by this field.

In this thesis, I augment previous behavioural models by studying their existence in environments with consumer learning. By extending static behavioural problems to dynamic environments with learning, I am able to explain puzzles in the areas of technology adoption and contract theory.

In chapter 1, I propose that status considerations – a feature of consumers' preferences overlooked by classical theory – can have positive effects in society whenever they are considered in an environment with active learning (i.e., experimentation). Status considerations - the fact that consumers care about how their wealth compares to their peers - have been consistently identified in the economic literature as a source of inefficiencies which should be both taxed and discouraged. Preferences for status are credited with distorting consumption decisions and with having a particularly harmful effect on developing countries by perpetuating poverty.

This paper suggests that status considerations can have a positive effect on experimentation, a crucial activity for innovation and technology adoption. Relying on recent empirical findings on properties of status preferences, I show that this effect exists whenever new technologies become available and have the potential to reduce the dispersion of final outcomes in society. Given the importance of technology adoption (driven by experimentation) for development, this result suggests that status considerations can have positive effects in a development context.

In chapter 2 and 3, I show that when naïve of behavioural consumers (who lack self-awareness about their preferences) can learn, pricing methods in subscription contracts, which were previously unexplained by standard contract theory, can be shown to be the optimal response of firms trying to prevent consumer learning.

In chapter 2, I propose a new mechanism to explain why firms in service sectors offer subscription contracts with three part tariff (3PT) pricing. Subscription contracts, where firms charge a fixed fee for access to the contract and then marginal prices for usage, are ubiquitous in service industries[2]. However, one of the most commonly used pricing schemes in these contracts - 3PT - cannot be explained by standard contract theory models. I am able to show that this pricing is indeed optimal by incorporating, into an otherwise standard contracting model, two empirically motivated behavioural assumptions: i) consumers are not fully aware of their demand for services, but can learn it over time; ii) consumers have imperfect memory, such that they recall past bills, but forget past usage.

I show that this imperfect memory implies that the firm can use its pricing today to influence future beliefs. Pricing will be used to manipulate the beliefs of consumers in the market, since consumer learning exacerbates the adverse selection problem of the firm. When consumer's overall valuation is high, the firm will face particularly strong incentives to manipulate the lowest beliefs in the market. Because consumers

[2]Examples are fitness (e.g., gyms); communications (e.g., mobile calls and data); utilities (e.g., electricity); banking (e.g., credit cards); online entertainment (e.g., Netflix); and so on.

with lowest beliefs consume less, the firm will offer 3PT: impeding consumer learning when consumers choose low usage and not when they choose high usage. This framework has novel implications for the regulation of cellphone markets.

In chapter 3, I augment the set-up of chapter 2 with three different extensions which establish the robustness of the original model to various considerations in this market and address potential weaknesses.

Chapter 1

Status and experimentation

Status, a non-standard component of consumers' preferences, incorporates into economic models the fact that consumers care about how their wealth is perceived by their peers. This concept is more than a theoretical curiosity. A large volume of empirical evidence (both in labs and markets) has consistently documented that Status is an important driver of consumers' choices. Despite its importance, there is no consensus on the economic literature about the welfare effects of preferences for status. Previous literature is divided into papers that emphasise status' negative welfare effects - arising from the negative externalities on zero sum status interactions - and status' positive indirect welfare effects - as status ability to correct other underlying problems in private or organisational incentives. Given the large potential welfare effects of status driven actions, previous authors have defended government intervention in both markets affected by status (which can potentially include luxury/status goods, real estate and labour markets) and in the educational content that can shape future generations' preferences for status.

This paper identifies a novel mechanism through which status can affect positively welfare in a dynamic information acquisition environment. I extend the classical experimentation framework (with bounded and public returns on technologies) by introducing empirically identified status preferences which imply concavity of status payoffs. I show that status is going generate a new incentive for information acquisition/experiment in an environment where, otherwise, there would be under-provision of information acquisition (due to free-riding). Concave status payoffs are going to increase incentives for experimenting on high potential unexplored technologies because coordination on these same technologies increases consumers' chance of ending up with same outcome and, hence, decrease status variance.

1.1 Introduction

Status, a non-standard component of consumers' preferences, incorporates into economic models the fact that consumers care about how their wealth is perceived by their peers. This concept is more than a theoretical curiosity. A large volume of empirical evidence (both in labs and markets) has consistently documented that Status is an important driver of consumers' choices. Despite its importance, there is no consensus in the economic literature about the welfare consequences of preferences for status.

To understand why the economic literature is divided about the impact of status, consider the simple case of status-driven consumption, known as "conspicuous consumption". This type of consumption, composed of expensive visible goods such as luxury goods, is motivated by wealth signalling with a view to achieving social status.[1]

Consumers' welfare – given by the sum of agents' intrinsic (consumption) utility and payoffs from status – can decline when status preferences lead to conspicuous consumption for two reasons. First, intrinsic utility decreases when status leads to signalling distortions on consumption choices (e.g. over-consumption of expensive sports cars or watches to signal wealth) which, in turn, can lead to negative consequences on such related decisions as savings and labour supply. Second, aggregate status payoffs resulting from conspicuous consumption are at best neutral. This follows from the zero-sum properties of status: when one consumer takes an action that leads him to higher status, another agent's status is decreased, since status is defined by a ranking. However, there can also exist indirect positive welfare effects. The status distortions exemplified above can correct incentives in otherwise distorted markets. For instance, over-consumption of expensive sports cars can generate higher incentives to work. This can lead to a positive welfare effect if in this example the labour market was suffering from an initial distortion (e.g. the existence of overly powerful unions) that would lead to a sub-optimal level of labour supply.

[1] Given status' social component, status driven decisions are strategic in nature and need to be analysed using game theoretical concepts of equilibrium.

Using these tools in the simple case where all consumers (who value status) can costlessly reveal their wealth, we can see that full wealth revelation is (in most frameworks) the unique equilibrium. By contradiction, an alternative (hypothetical) equilibrium where no one reveals their wealth would not be stable: every consumer with wealth above average would deviate by revealing to be wealthier than the average. Recursively, this argument would unravel any equilibrium without full wealth revelation.

In more realistic cases where wealth cannot be revealed easily and credibly to peers, we need to consider signalling models used to analyse status decisions in equilibrium.

This example underscores the importance of understanding the mechanisms through which status affects welfare. Even in this simple case, these effects are ambiguous and depend on the relative size of the costs from signalling, versus the indirect (potential) benefits from having a separating equilibrium with status. Furthermore, this discussion has important policy implications, because status can be influenced by policy choices such as education (e.g. should children be taught to be the best in their class or simply aim to be successful in absolute terms?); and taxes (e.g. should consumption driven by status be penalised or incentivised?).[2]

The model proposed here aims to improve our understanding of the mechanisms through which status can positively affect welfare. In particular, when status follows empirically-based assumptions, I show that equilibrium decisions on innovation and technology adoption can be improved by the existence of status preferences.

To study this problem I consider the classical model of experimentation with public information, a framework developed to analyse dynamic decisions of information acquisition. In the standard version of this model, status considerations are absent. A known result in this set-up is that, because information is public, agents free-ride on costly information acquisition and equilibrium information acquisition is sub-optimal. On the basis of this result, the model has been used to rationalise such issues as lack of innovation and the under-adoption of technology (e.g. vaccination and water filters in the context of development).

I augment this framework by assuming consumers' preferences include status considerations and show that such consumers invest more in information acquisition than their counterfactual "standard consumers" (who ignore status in their decision-making).

The model. The classical experimentation framework is used to develop the basic set-up. Over several periods, homogeneous risk-neutral consumers face a choice between status quo technologies with known returns and new technologies with unknown returns. Data from the outcomes of new technologies can be used to learn about them and, hence, improve future decisions. This feature makes this an experimentation model: present decisions between new and status quo technologies are driven by the future value of acquiring more information about actions - i.e.

[2]See Heffetz and Frank (2008) for a discussion of policies designed to curb expenditure on "status goods" using taxes on luury goods and progressive consumption taxes. This same paper argues that recent inequality growth makes this discussion particularly urgent as it magnifies problems caused by status. This prediction is confirmed by signalling models and it is relevant regardless of whether inequality trends are real or just perceived (see Autor, Katz and Kearney (2008) for recent measurement issues on inequality).

"experiment". [3]

I consider the simplest form of the model:[4] the decisions are made over two periods and the returns to technologies, driven by a simple parameter (probability of success), are binary and labelled 'success' and 'failure'. The agents' choice is also reduced to two technologies: a single status quo technology - safe option - with known but moderate probability of success; and a single new technology – a risky option. The risky option's probability of success is fixed over time but unknown to the agents initially, and can be either high or low.

The non-standard element of this set-up is that the agents' utility is the sum of an intrinsic utility component – determined by the absolute binary returns from their chosen technology – and a status payoff component – determined by the comparison between their absolute returns and their opponent/peer. Inspired by recent empirical observations, finally, status payoffs are assumed to be concave such that they increase at a decreasing rate.

Throughout the paper, I focus on a simple development application that both focuses the terminology and motivates some assumptions. I assume that the technology at stake is a water filter which determines the outcome of each agent: being healthy or sick (in each period). The two available technologies are a traditional water filter with a known success rate, and a new water filter with an unknown success rate (which can either be high or low). The use of these techniques motivates the technology structure assumed. In particular, the preventive health techniques here do not provide long-term immunity (as vaccination) and do not affect the transmission rate to or from other people (as HIV). For simplicity, the main outcome of the preventive health technique (the health status) correlates perfectly with income, such that healthy people are more wealthy than sick people. While this example is helpful in making the problem at hand more recognisable, the results for this particular application should note be taken literally as there are many other important dimensions relevant for technology adoption problems in developing countries.

Results. In a preliminary result, I establish a benchmark without status by considering the case where status payoffs are negligible, implying that agents' decisions are driven solely by their intrinsic utility. In this case where status motivations are absent, a standard result (shared with previous experimentation literature) es-

[3]Hence, consumers in these models are commonly said to choose between the "exploitation" of known choices and the "exploration" of unknown choices (that is, to experiment).

[4]Without loss of generality, this framework could be generalised for more general technologies and choices. See the discussion at the end of the chapter for a more involved treatment

tablishes that, in equilibrium, information acquisition/experimentation is inefficient. This result follows from the public nature of information: consumers are going to under-experiment because they internalise fully their costs of information acquisition but not the overall (public) benefits. Using this case as a benchmark, the main result considers then equilibrium experimentation when status payoffs are large. In this case, under some assumptions, status is going to increase the incentives to experiment and, hence, increase welfare. This result relies on two observations:

First, status dispersion is lower when both agents choose the new (risky) choice. New technologies have the potential to be highly successful or highly unsuccessful. This implies that when both agents coordinate on the same new technology they are much more likely to get the same outcome, that is, either both are successful or both unsuccessful. This means that coordinating on the new technology leads to much less outcome variation between the agents and, hence, less status variation. This property can be seen in the application used throughout the paper. Suppose that the success rate of a water filtering technology is the probability of a person's being healthy after using it. Assume that the traditional water filter has a known success rate of 50%, while the new water filter has an unknown success rate: it is either good (90% success rate) or bad (10% success rate). Then, ignoring contagion, it is easy to see that the probability of agents having different outcomes (i.e. one being sick and the other is healthy), when both use the same technique, is 50% if they use the traditional technique and 18% if they use the new one.

Second, consumers (with concave status payoffs) dislike status dispersion. The assumption of concave status payoffs - based on implications from recent empirical findings on the consumption of visible goods - implies that the benefits from status (in terms of utility) are decreasing. This assumption implies that consumers dislike status variation. The argument for this fact follows a similar textbook intuition that implies that consumers with concave utility in wealth are risk averse. Putting these two facts together, it is easy to show that when consumers have concave status returns they will have higher incentives to choose new risky options which lead to less status dispersion.

Main contribution and caveats. First, this model suggests a new mechanism through which status positively affects the welfare of economies. This paper shares with the existing literature the argument that status can increase welfare by correcting distortions that would exist in a world without status. However, the mechanism proposed here, contrary to the previous literature, does not rely on the fact that status increases individual returns to success, when one consumer is able to get ahead of its peers. Instead, in this model, individual returns to success (in experimentation) are lower in two ways.[5]First, status returns are assumed to be concave, implying that experimentation leads to smaller individual status gains and greater individual status losses. Second, public information implies that the future gains from information are shared, leading to zero individual status gains for the agents who experiment.

The novelty of this mechanism relies on the fact that status increases the benefits of coordinating on more risky options. This effect is driven by concave status payoffs, an empirically-based assumption, which changes the incentives by making the status quo options – with relatively wide status variation - less attractive than a new option with smaller status variation. Moreover, the main result of this paper can also be used to discuss optimal incentive systems in firms. Many large mature tech firms in rapidly changing sectors consider the problem of staying competitive by continuing to innovate. What is the optimal way for these firms to use their substantial resources to create an innovative culture in their company? Can relative performance schemes be the answer to this question. The present paper suggests that firms that want to foster innovation and risk-taking beyond the level that individual incentives can generate should use "concave" performance bonuses based on relative performance across peers. This system, which mimics the status payoffs discussed in this paper, can lead to higher incentives for experimentation when the firm faces an environment for information acquisition in which its ventures/projects' outcomes are bounded and stochastic. The effect of status payoffs developed here mirrors similar results developed on a large literature on "tournaments" - optimal relative performance schemes[6]. This literature does not focus on information acquisition decisions but has shown similar qualitative results for optimal effort incentive provision. For instance, when tasks considered require coordination between employers participating in the tournament, excessively high incentives to success might lead to sub-optimal effort levels[7].

Finally, it is important to discuss some of the assumptions on which this paper's results rely. By focusing on the experimentation problem, this model ignores other costs of status discussed above such as those driven by signalling wealth through the

[5]In this set-up, each consumer faces an investment decision in information acquisition with a typical payoff structure: lower expected outcomes today can lead to (potentially) higher future outcomes.

[6]Introduced by Lazear and Rosen (1981).

[7]As reviewed in McLaughlin, 1988

consumption of visible goods. This is done to keep the analysis simpler and implies that our analysis cannot address the question of which effect will dominate in the overall welfare analysis. Similarly, to make inference decisions simple, it is assumed that agents perfectly observe information from other agents. In the last section, it is shown that this assumption could easily be relaxed without changing the results so long as communication between agents is allowed.

1.2 Related literature.

Status. The literature on status is supported by considerable empirical evidence, surveyed by Frank and Heffetz (2008), that status has an important impact on consumers' decision making. There is empirical evidence on reported preferences. On world surveys, it is a well-known fact that consumers self-reported well-being is more strongly correlated to relative than to absolute income[8]. Similar evidence, found on experimental designs, suggests that people prefer individually worst options when these imply a higher relative ranking when compared to their peers, as shown by Ferrer-i-Carbonell (2005) and Duesenberry (1949).[9] Moreover, there is direct evidence in consumption markets that consumers are willing to spend considerable proportion of their income on highly visible consumption (also called conspicuous consumption) to signal their wealth. In particular, Charles, Hurst and Roussanov (2008) show that the consumption of highly visible goods in the US has very different characteristics from "standard" consumption and, in particular, it is going to depend on the income of consumers' peer groups as theoretical models of status would propose.

The assumptions on the properties of status payoffs used in this paper are motivated by findings of this empirical literature about the importance of status for consumers over different income levels. Results from the large data set studied by Charles, Hurst and Roussanov (2008) suggest, using a signalling model to explain consumption evidence in the US, that the perceived utility gains from status are decreasing as the level of wealth increases. This model is used to explain evidence that poorer people spend more in relative terms on visible/status goods than richer people. These conclusions are not contradicted by anecdotal and empirical observations in developing countries. For instance, in North India, the average proportion of income spent on weddings, which accounts for a third of the household's annual income for poorer households, is decreasing with income level (Bloch et al (2003))[10].

This paper contributes to a theoretical literature which studies the welfare effects of preferences for status in markets. This literature, discussed below, can be split into models that predict positive and negative welfare effects arising from status preferences.

The negative welfare effects of status have proven to be relevant in a variety

[8]See, for instance, Veenhoven (1993) although this evidence is not without controversy as argued by Stevenson and Wolfers (2008)

[9]Similar results have been found on job satisfaction reports in Clark and Oswald (1996).

[10]And this fact is not explained by indivisibilities.

of classical theoretical frameworks: decision theory (Robson and Ray (2012)); general equilibrium (Arrow and Dasgupta (2009)); evolutionary economics (Samuelson (2004)); development (Moav and Neeman (2010))[11], etc. In most of these papers, status driven actions lead to problems akin to traditional externalities. Externalities arise from the fact that any action that increases an agent's relative status is going automatically translate into a relative decrease to (some of) the remaining agent's relative status, by definition of the status ranking. This mechanism has led some authors to compare the inefficient equilibrium in status games to the one in a Prisoner's Dilemma (Congleton, 1989). Importantly, supply side competition (on firms' supplying goods used to signal status) does not seem to attenuate this problem as discussed in detail in Bagwell and Bernheim (1996) and Becker, Murphy, and Glaeser (2000). The conclusions of most of these papers, imply a role for policy makers in improving markets using among other tools taxation of status goods as argued by Frank (1985) and Ireland (1994).

On the other hand, there is a smaller literature which argues that status can have positive welfare effects in society. Arrow (1971) and Weiss and Fershtman (1998) argue that status can have a role as a non-monetary/cultural incentive system that allows societies to create incentives for actions that markets fail to reward appropriately. In the same spirit, Glaeser et al. (2000) argue that status can be used as a proxy for important social capital and intangible people's skills. That is, status can have a positive effect if it can be used to decrease important asymmetric information problems.

Technology adoption and experimentation. Experimentation literature in economics, also known as "bandit problems" and surveyed by Valimaki and Bergeman (2006), was started by Rothschild (1974). This literature provides a stylised framework to analyse information acquisition in dynamic environments. In this "active learning" models, agents choose repeatedly between actions which have known returns and actions which have (initially) unknown returns. The solution to this problem requires balancing the classic trade-off between exploitation (of known actions) and exploration (of uncertain actions with outcomes which can be higher or lower than the known actions).

When more than one agent experiments technologies with information spillovers, this decision becomes strategic and, as shown by Bolton and Harris (1999), there can

[11] Moav and Neeman (2010) formalise the intuition that small mistakes driven by status are particularly costly in poverty environments, showing that status can create poverty traps.

be free-riding leading to sub-optimal levels of experimentation.[12] This is a similar argument to the one developed in Besley and Case (1994)[13] that, in decentralised economies, experimentation is likely to be under-provided because of the public good nature of information. They argue that institutions may be needed to internalize public gains of information when free-riding problems are likely to be an issue.

[12]Other effects also studied by these authors can change these conclusions, such as the encouragement effect. However, the encouragement effect is not as general as the free-riding effect as shown by Keller and Sven (2010). They show that the strength of each of these effects depends crucially on the properties of the stochastic processes that drive the actions and the encouragement effect is not robust to many of the classical specifications in these environments.

[13]References for the role of public information on under-experimentation can also be found in Foster and Rosenzweig (1995)

1.3 Set-up

Typical experimentation set-up. There are two agents $i \in \{1, 2\}$ with symmetric preferences who live for two periods and maximise discounted payoffs. In each period t, agents choose, simultaneously and independently, a preventive health technique $x_{i,t}$ for $i = 1, 2$. There are two health techniques, labelled risky (R) and a safe (S), which are fully characterised by an effectiveness parameter $\theta_x \in (0, 1)$. The effectiveness of each technique is fixed over time and common to all agents. However, the consumers are not fully informed about both techniques: while θ_S is known, θ_R is unknown. Both agents share a common prior about technique R: it is either "good" $\theta_R = \bar{\theta}_R$ with probability α_1; or "bad" $\theta_R = \underline{\theta}_R$ with probability $1 - \alpha_1$. The techniques' effectiveness satisfy the following relationship: $\underline{\theta}_R < \theta_S < \bar{\theta}_R$. While not crucial for the results, it simplifies the analysis to assume symmetry around 1/2:

$$\theta_S = 1/2 \qquad\qquad | \, 1/2 - \underline{\theta}_R \, | = | \, \bar{\theta}_R - 1/2 \, | = \Delta < 1/2$$

Payoffs. Consumer i's per period utility $u_{i,t}$ is given by the addition of intrinsic payoff and status payoffs:

$$u_{i,t} = m_{i,t} + \zeta \, l(m_{i,t} \, , \, m_{-i,t} \,)$$

Where intrinsic payoffs are determined by the agent's income $m_{i,t}$; and the status payoffs are determined by $l(m_i \, , \, m_{-i} \,)$, which depend on the income of the consumer as compared to his peer. The importance of status payoffs in utility is captured by the a fixed weight $\zeta \in [0, \infty)$, such that these are standard consumers[14] when $\zeta = 0$. Status payoffs follow a particularly simple form given by:

$$l(m_i, m_{-i}) = \begin{cases} g & \text{if} \quad m_i > m_{-i}; \\ 0 & \text{if} \quad m_i = m_{-i}; \\ -1 & \text{if} \quad m_i < m_{-i} \end{cases}$$

For simplicity, assume that consumers' income is determined by their health status, such that $m_{i,t} = H_{i,t}$ which allows us to rewrite utility per period as:

$$u_{i,t} = m_{i,t} + \zeta \, l(m_{i,t} \, , \, m_{-i,t} \,) = H_{i,t} + \zeta \, l(H_{i,t} \, , \, H_{-i,t} \,)$$

Which implies that the expected status payoffs are given by:

$$L(x_i, x_{-i}) \equiv E[\, l \,(H_i, H_{-i}) \,|\, x \,] = Prob(H_i > H_{-i} \,|\, \mathbf{x} \,) \, g - Prob(H_i < H_{-i} \,|\, \mathbf{x} \,)$$

[14] Without status considerations.

Therefore, the overall expected payoffs in each period are determined by the expected probabilities of being: healthy (for intrinsic payoffs); and the healthiest agent (for status payoffs).

Health status and preventive health techniques. The probability of an agent i being healthy is determined by choices in the current period and depends on both by the effectiveness of the technique chosen by the agent $x_{i,t}$, and by the contagion rate from the other agent $C(\theta_{x_{-i}})$. We use a simple stylised model of contagion where the agents' final health status H is determined over two sequential stages: infection and contagion. First, each consumer (who chose x) has contact with the disease vector which determines the agent's intermediate health status h. During this stage, the agent either stays healthy (i.e., $h = 1$) with probability θ_x, or gets sick ($h = 0$) with probability $1 - \theta_x$. Second, an agent i can, with probability $\gamma \in [0, 1]$, get the disease transmitted from another agent. This can only happen when the other agent was infected in the first stage, that is $h_{-i} = 0$, implying that the probability of contagion from agent $-i$ is given by:

$$C(\theta_{x_{-i}}) \;\; = \;\; \gamma\,(1 - \theta_{x_{-i}})\,)$$

Hence, the probability that agent i is healthy (i.e., $H = 1$) in the end of each period depends on the actions chosen by both agents and it is given by:

$$P_i(\mathbf{x}) \;\; = \;\; P(x_i, x_{-i}) \;\; = \;\; \theta_{x_i}\,(\,1 \;-\; C(\theta_{x_{-i}})\,)$$

Similarly, the probability that agent i has high status (i.e., agent i is healthy and agent $-i$ is sick but does not transmits disease) is given by:

$$Prob(H_i > H_{-i} \mid \mathbf{x}\,) \; = \; E[\,\theta_{x_i}\,(\,1 - \theta_{x_{-i}} - C(\theta_{x_{-i}})\,)\,]$$

Overall decision and basic assumptions. The timing of each period t is simple: agents start each period healthy, choose a technique $x_{i,t}$ and the vector of techniques $\mathbf{x_t} = (x_{i,t}, x_{-i,t})$ determines their final period specific health status $H_{i,t} \in \{0, 1\}$ and, hence, determines each agent's payoffs in period t.

Under the following two assumptions and as long as agents do not care about status (i.e., $\zeta = 0$), the set-up described becomes a classical experimentation model with public information. First, in order to consider an "experimentation" environment, rule out the cases where information acquisition decision is a secondary problem (that is, cases where choosing R is always optimal):

Assumption 1: ***Pessimism and Short-sightness.***

$$\alpha_1 < 1/2 - \frac{\gamma\Delta}{2-\gamma} \qquad \text{and} \qquad \delta < \frac{1-\gamma/2}{1/4-\Delta^2}$$

In other words, the risky action is unlikely to be in a good state and the agents are sufficiently impatient. Under this assumption, there is a present cost of choosing R, implying that agents only choose R if future benefits (of information acquisition) are high enough.[15]

Second, assume that agents learn rationally and that there is full observability of previous outcomes and actions, making this a public information environment:

Assumption 2: *Information about experimentation is public. At the beginning of each period t, each agent i observes all past information from the other agent, that is observes the technique choice $x_{-i,t-1}$ and intermediate health status $h_{-i,t-1}$.*

This assumption on full observability can be relaxed as analysed in the discussion section. Under this assumption, agent's have common beliefs in period 2 given by α_2. These beliefs are determined by Bayes rule using first period's observations: when there is no experimentation in the first period (i.e., $x_1 = S, S$), there is no learning and $\alpha_2 = \alpha_1$; when there is experimentation in the first period by one agent (that is, $\mathbf{x_1} = R, S$ or $\mathbf{x_1} = S, R$) the posteriors depend on information observed as follows:

$$\alpha_2 = \begin{cases} \alpha^+ = \frac{\alpha\bar{\theta}_R}{\alpha\bar{\theta}_R + (1-\alpha)\underline{\theta}_R} & \text{if observed } h_R^i = 1; \\ \alpha^- = \frac{\alpha(1-\bar{\theta}_R)}{\alpha(1-\bar{\theta}_R)+(1-\alpha)(1-\underline{\theta}_R)} & \text{if observed } h_R^i = 0. \end{cases}$$

And when both agents experiment in the first period (that is, $x_1 = R, R$), the posteriors are:

$$\alpha_2 = \begin{cases} \alpha^{++} = \frac{\alpha(\bar{\theta}_R)^2}{\alpha(\bar{\theta}_R)^2+(1-\alpha)(\underline{\theta}_R)^2} & \text{if observed } h_R^i = 1\ ,\ h_R^{-i} = 1; \\ \alpha^{+-} = \alpha_1 & \text{if observed } h_R^i = 1\ ,\ h_R^{-i} = 0; \\ \alpha^{--} = \frac{\alpha(1-\bar{\theta}_R)^2}{\alpha(1-\bar{\theta}_R)^2+(1-\alpha)(1-\underline{\theta}_R)^2} & \text{if observed } h_R^i = 0\ ,\ h_R^{-i} = 0. \end{cases}$$

Finally, we augment this experimentation model by introducing empirically motivated assumptions on status discussed on the introduction, such that status payoffs are positive and concave:

Assumption 3: ***Concave status.*** $\qquad 0 < g < 1 \quad \text{and} \quad \zeta > 0$

[15] This assumption allows us to focus on the experimentation decision by ruling out cases where: there is no present cost of choosing R; and the weight on future benefits is so large that there is no meaningful trade-off.

1.4 Preliminary results (without status)

Start by considering the problem when $\zeta = 0$ that is the consumers do not care about status and decisions are driven by intrinsic utility.

First period. From the point of view of the first period, the payoffs of agent i when ignoring status payoffs, consists of optimising (given the expected actions of the other player) are:

$$P(\, x_{i,t}, x_{-i,t}\,) \;\; + \;\; \delta \, E[\, P(\, x_{i,t+1}, x_{-i,t+1}\,) \,|\, x_t\,]$$

Then, under assumption 1 on α_1, choosing R in the first period decreases the probability of being healthy:

$$P(R, S) \;\; < \;\; P(S, S) \qquad \text{and} \qquad P(R, R) \;\; < \;\; P(S, R) \qquad \text{at}$$

$\alpha = \alpha_1$

Hence, agents will choose R in the first period only when future expected benefits from information are high enough:

Lemma 1. *Under assumption 1, choosing R in the first period is experimenting. That is, choosing the risky (R) option is costly in terms of first period payoffs and it will only be done if future informational benefits are large enough.*

To understand the experimentation decision, we need to characterise the second period's benefits of experimentation.

Second period. The benefits of experimentation depend on the relationship between second period's equilibrium and the agents' beliefs, given by:

Lemma 2. *(Pure strategy) equilibria in the second period are given by:*

- S, S *when* $\alpha_2 \;\leq\; 1/2 - \frac{\gamma \Delta}{2+\gamma}$
- R, R *or* S, S *when* $\alpha_2 \;\in (\, 1/2 - \frac{\gamma \Delta}{2+\gamma} \,,\, 1/2\,)$
- R, R *when* $\alpha_2 \;\geq\; 1/2$

Agents coordinate on the same action, because there is a complementarity when choosing R:

$$P(R, R) - \; P(S, R) \quad > \quad P(R, S) - \; P(S, S) \qquad \text{for any}$$

α

When agents coordinate on R, the contagion decreases compared to when only one chooses R. This effect exists because contagion only occurs when agents have different intermediate health status (i.e., $h_i \neq h_{-i}$), which happens with probability $\theta_{x_i}(1-\theta_{x_{-i}})\gamma$. By coordinating on R, they are both either more likely to be healthy or sick (than when choosing S), making contagion less relevant. This effect can be seen below in the graph below which represents the probability of contagion when agents coordinate on an action with effectiveness θ:

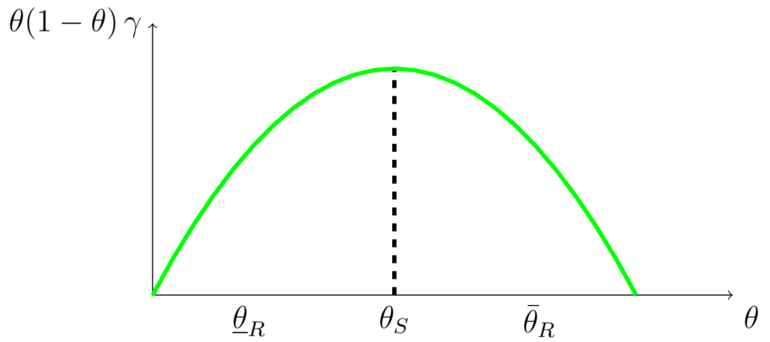

Figure 1.1: Probability of contagion when coordinating.

When second period beliefs are given by α_2, the second period (myopic) decision of agents leads to the equilibrium above, determined solely by the relationship between probabilities of being healthy. Given this second period equilibrium, we can characterise the benefits of experimentation in the second period:

Lemma 3. *The second period's benefits of experimentation are:*[16]

i) Positive for experimenting with a single agent when $\alpha_1 > 1/2 - \Delta\frac{1+\gamma/2}{1+2\Delta^2\gamma}$

ii) Higher when two agents experiment, instead of one.

The flexibility of being able to adjust actions to the signals observed in the first period creates an option value benefit in the second period. In particular, after the consumer experiments it can switch back to S if it gets a negative signal about R. The option value benefit relies on the fact that the information observed can potentially change the agent's actions, hence experimenting only brings expected benefits in the second period if R is likely enough to be in a good state[17].

[16]Note in this lemma the net future gains from experimentation computed are: the difference in second period's payoffs that follow $x_1 = R, S$ and $x_1 = S, S$ in i); and the difference in second period's payoffs that follow $x_1 = R, R$ and $x_1 = R, S$ in ii).

[17]Otherwise, even a positive signal will not move consumers' beliefs enough to change their action.

Under-experimentation: equilibrium and efficiency. Throughout this section, we ignore technical issues with multiple equilibrium by assuming that whenever there is more than one Nash equilibrium the agents choose to coordinate on the one where there is experimentation. This is without loss of generality, as obviously relaxing this equilibrium selection makes the under-experimentation result easier to prove. To characterise experimentation in equilibrium we are particularly interested in the experimentation frontier, that is the threshold in beliefs from which in equilibrium (as defined above) at least one agent experiments:

Lemma 4. *When x_2 follows lemma 2 above*[18]*, experimentation frontier is characterised by:*

- *Pioneering equilibrium:* $x_1 = R, S$, *whenever* $\alpha_1 > \alpha^* \equiv 1/2 - \frac{\delta\,\Delta\,(1+\gamma/2)}{2\,(1-\gamma/2+\delta/2\,(1+2\,\gamma\,\Delta^2))}$ *and* $\delta \geq \hat{\delta}$.
- *Joint experimentation equilibrium:* $x_1 = R, R$, *whenever* $\alpha_1 > \alpha_J^* \equiv 1/2 - \frac{\Delta^2\,\gamma\,(1-(1/4-\Delta^2)\,\delta)}{2-\gamma-\delta\,2\,(1/4-\Delta^2))}$ *and* $\delta < \hat{\delta}$.

Where $\hat{\delta}$ is defined by comparison of the two thresholds.

This lemma follows from two facts. First, the first period's costs of experimentation are lower when both consumers experiment by complementarity of R, meaning that if consumers don't care much about the future (i.e., δ is small), then the experimentation frontier is defined by the equilibrium with two agents experimenting. Second, (for the parameters considered) the benefits of experimentation - in terms of second period's payoffs - are higher for a pioneer than for an agent joining experimentation, as the benefits of more information are decreasing. This means that as agents care more about the future (i.e., δ increases), the experimentation frontier will be given by the equilibrium with one agent experimenting.

To develop an efficient benchmark consider a central planner (CP) who maximises the sum of the utility (without status) of both agents and experiments when $\alpha_1 > \alpha^{CP}$ with one agent and when $\alpha_1 > \alpha_J^{CP}$ with two agents. When comparing the sub-game perfect equilibrium with this efficient benchmark, note that the CP takes into account two factors that the individual agents disregard:

i) Negative externalities from contagion. That is the probability of being healthy decreases when the other agents experiments: $P(S,S) > P(S,R)$ (for any α_1). While

[18]In other words, in a subgame perfect equilibrium with equilibrium refinement defined above.

agents do not take this into account the CP does, leading to too much experimentation.

ii) Positive externalities from public information[19]*, as experimentation by one agent provides useful information for both agents.* The central planner values the future twice as much as the agent that is deciding to experiment, since it values the benefits accrued by both agents. It is easy to show that equilibrium experimentation increases (i.e., α^* decreases) whenever they value the future more (i.e., δ is higher), meaning that this factor makes agents under-experiment.

These two factors drive an wedge between agent's experimentation and the efficient benchmark. While the first factor means that in equilibrium there can be too much experimentation, the second means that there can too little experimentation. In the end, there is under-experimentation (when comparing equilibrium experimentation to the CP's benchmark) because contagion is a second order effect compared to the direct benefits of experimentation:

Proposition 1. *There is under-experimentation in equilibrium: CP experiments for a wider range of technologies (as characterised by their likelihood of success/being in a good state α_1, such that*

- $\alpha^* > \alpha^{CP}$*, when one agent experiments at the experimentation frontier.*
- $\alpha_J^* > \alpha_J^{CP}$*, when two agents experiment at the experimentation frontier.*

[19] As techniques have common properties for both agents.

1.5 Results (with status)

To focus on how status payoffs drive decisions, consider the problem when ζ is large enough, so we can disregard intrinsic utility, i.e., consumers choose as if caring only about status payoffs.

First period. The overall status payoffs for player i as evaluated from the first period are given by:

$$L(\, x_{i,t}\, ,\, x_{-i,t}\,) \;+\; \delta\;\; E[\, L(\, x_{i,t+1}\, ,\, x_{-i,t+1}\,)\, |\, x_t\,]$$

Techniques chosen affect expected status by changing the probabilities of having higher, lower or the same income/health status. It is useful to adapt assumption 1:

Assumption 1′: ***Pessimistic priors.*** $\alpha_1 < 1/2 - max\{\, \Delta\, \frac{1-g}{1+g}\, ,\, \Delta\, \frac{\gamma}{2-\gamma}\, \}$

This variation is equivalent to the previously defined assumption 1 whenever: $g \geq 1 - \gamma$. Under this assumption, choosing R in the first period decreases the probability of being the healthiest and increases the probability of being the least healthy and, hence, decreases expected status:

$$L(R, S) \;<\; L(S, S) \qquad \text{and} \qquad L(R, R) \;<\; L(S, R)$$

This allows us to develop a result (similar to the one on intrinsic utility) on first period's payoffs:

Lemma 5. *Under assumption 1', choosing R in the first period is similar to "experimenting" in terms of status payoffs: it is costly in the first period and can potentially bring status benefits in the second period.*[20]

Implying that, as before, agents only choose R/experiment in the first period when future benefits are high enough to motivate them to choose the risky technique.

Second period. As before, choosing R is more advantageous when the other agent also chooses it:

Lemma 6. *There are complementarities in R:*

$$L(\, R, S\,) \;-\; L(\, S, S\,) \;<\; L(\, R, R\,) \;-\; L(\, S, R\,)$$

[20]Alternatively, condition can be written as a new assumption 1: $\alpha_1 \;<\; 1/2 - \frac{\Delta(1-g)}{g+1}$.

As before, coordinating on R decreases the probability of having different health outcomes, as when agents choose the same action the probability of having different health status is given by $\theta(1-\theta)(1-\gamma)$. Therefore, when both agents choose the same action $x_i = x_{-i}$, the probability of having the same health status (*drawing the status game*) can be represented as:

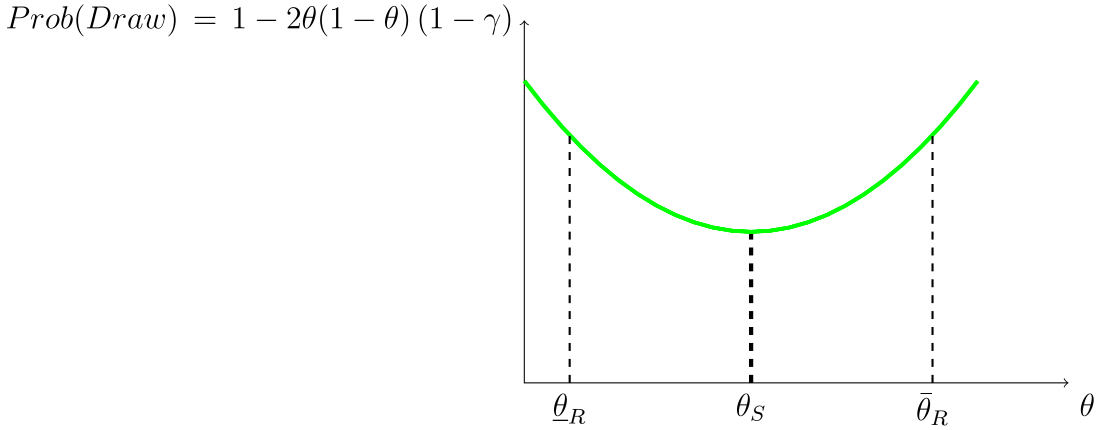

Figure 1.2: Probability of drawing when coordinating.

Implying that more "extreme techniques" lead to status draws more often. As status disparity decreases payoffs (since $g < 1$), more drawing is a positive effect.

There are two implications of this property:

First, a similar coordination equilibrium is generated in the second period when status is taken into account:

Lemma 7. *(Pure strategy) equilibria in the second period (with large ζ) is given by:*

- S, S *when* $\alpha_2 < 1/2 - \frac{\Delta(1-g)}{g+1}$
- R, R *or* S, S *when* $\alpha_2 \in (1/2 - \frac{\Delta(1-g)}{g+1}, 1/2)$
- R, R *when* $\alpha_2 > 1/2$

Under assumption 1' and using the second period's posteriors described in the set-up, if there is no experimentation in the first period the equilibrium in the second period is S, S (as $\alpha_2 = \alpha_1$). If there is experimentation, the resulting equilibrium in the second period depends on the information observed.

Second, in terms of future payoffs, there is no incentive to *join experimentation when the other agent is experimenting* (that is, deviate from $x_1 = S, R$ to $x_1 = R, R$ equilibrium):

Lemma 8. *Under assumption 1',*[21] *the second period's status net benefit from joining experimentation are negative.*

First, note that beliefs do not play a role on the second period's payoffs for each particular equilibrium. To understand this point note that, by symmetry, the payoffs in R, R equilibrium do not depend on beliefs (such that $E[\,L(R,R)\,|\,h = 1\,] \;=\; E[\,L(R,R)\,|\,h = 1, h = 1\,] \;=\; L(R,R)$) and are given by:

$$L(R,R) \;=\; E\,[\,\theta_R(1-\theta_R)\,(1-\gamma)\,] \;\;=\;\; 1/4\,(1-\gamma)\,(g-1) \;+\; \Delta^2\;(1-\gamma)\,(1-g)$$

Which, given that $g < 1$ (assumption 3), these status payoffs are larger than when both agents choose S action:

$$L(S,S) \;=\; \theta_S(1-\theta_S)\,(1-\gamma)) \;=\; 1/4\,(1-\gamma)\,(g-1) \;\;<\;\; L(R,R)$$

So the gains from experimentation in terms of status payoffs come mainly from coordinating on an action with less status dispersion.

Second, agents are less likely to coordinate on R, R following a first period where both agents experimented (i.e., $x_1 = R, R$) than when only one experimented (i.e., $x_1 = R, S$). This follows from the fact that getting two positive signals (i.e., $h_R = 1$) is less likely than getting one:

$$prob(h_R = 1) \;>\; prob(h_R = 1, h_R = 1)$$

And beliefs following experimentation are given by the probability tree[22] below:

Where, importantly, the posteriors after observing one positive and one negative signal are equal to the prior:

$$\alpha^{+-} \;\equiv\; prob(\bar{\theta}_R\,|\,h_R = 1, h_R = 0) \;\;=\;\; \alpha_1$$

Meaning that the equilibrium in this case is S, S.

[21] And with similar refinement on multiple equilibrium as defined in the previous section

[22] Using beliefs defined in the set-up.

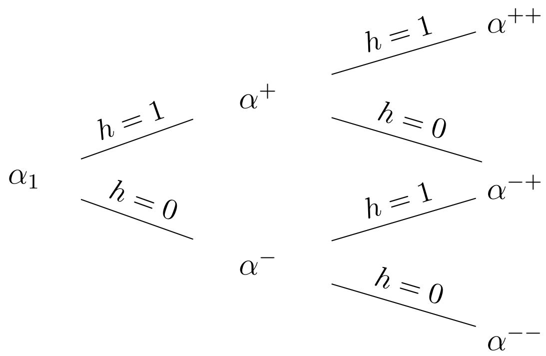

Figure 1.3: Probability tree following experimentation.

Therefore, since future status benefits are maxisimised with one single agent experimenting, agents will never join experimentation since experimenting is costly.

A further implication of the fact that beliefs play a small role in future status benefits of experimentation is that:

Lemma 9. *For any $\zeta > 0$, the second period's benefits of experimentation from status are weakly larger than the benefits from intrinsic utility, whenever α_1 is small enough.*

In intuitive terms, gains from intrinsic utility payoffs come from the possibility of learning positive information, which influences the chances of getting sick in the next period. These gains are given by:

$$prob(h=1)\ (\ E[\,P(R,R)\,|\,h=1\,]\ -\ P(S,S)\ ) \qquad \text{for } \alpha_1 > 1/2 - \Delta$$

Gains from status come from the fact that the agents coordinate on a technology which leads to smaller probability of different outcomes - decreasing status dispersion.

$$prob(h=1)\ (\ L(R,R)\ -\ L(S,S)\ ) \qquad \text{for } \alpha_1 > 1/2 - \Delta$$

Hence, the second period gains from experimentation from intrinsic utility are more strongly related to probabilities of good state α_1 and will decrease faster in this parameter.

Proposition 2. *Whern consumer's preferences are mainly driven by status payoffs* *(i.e., for large enough ζ), consumers experimentation in equilibrium is efficient if they value the future enough such that $\delta = \delta^S$. In this case, the experimentation frontier coincides with CP decision to experiment, where:*

$$\delta^S \equiv \frac{1 + \gamma + 3g + g\,\gamma}{(1 - 4\Delta^2)\,(1 - g)}$$

The experimentation frontier of both CP and status driven agents (in decentralised equilibrium) is going to be positively related to the discount factor δ, that is the lower bound on α_1 decreases as δ increases. This follows from the simple fact that experimentation is driven by future benefits that are more valuable as the agents care more about the future.

When discount factor is large enough, experimentation frontier happens when the expected level of success of R is low (i.e., α_1 is low). At that level, by lemma 9, the experimentation gains of status are higher because status gains do not come mainly from a coordination benefit that decreases dispersion of outcomes, while benefits from intrinsic utility are more strongly related to the likelihood of R being in a good state (that is, α_1).

1.6 Discussion.

The results of this paper rely on the fact that, concave returns from status imply that there is a net gain when agents experiment with a technology which has the potential to decrease status dispersion, as agents give up two equally likely events: gains of $g \in (0, 1)$ and loses -1 (for zero payoffs).

The results in this paper are achieved under a few simplifying assumptions discussed here. First, symmetry of probability distributions of technology R plays an important role in simplifying the analysis of this problem, but could be easily relaxed. While, this would mean that gains from status would be related to the probability of each state of R, this relationship would be of second order importance under standard regularity conditions, and results would be qualitatively similar. Second, agents are assumed to observe intermediate health status h, instead of the final status. If instead agents observed only final health status, the problem would not change when both agents experiment, but it would change if only one agent experiments because there would be an important divergence of consumer's beliefs. However, allowing consumers to communicate their private information would solve this problem since consumers would always want reveal their intermediate health status whenever g is small enough. This is a consequence of the status benefits from experimentation coming from coordinating on R. Finally, this model could be extended to more periods, but the analysis would have to take into account additional effects (as the encouragement effect discussed in Bolton and Harris, 1999) which could mitigate the free-riding effect in benchmark model and complicate significantly the analysis.

1.7 Appendix

1.7.1 Appendix of section 1.4

Proof of lemma 1 and lemma 2. These lemmas are immediate using the properties of probability of being healthy defined as:

$$Prob(\,H_i = 1\,|\,\mathbf{x}\,) \;\equiv\; P(x_i, x_{-i}) \;=\; E[\,\theta_{x_i}\,[\,1-(1-\theta_{x_{-i}})\,\gamma\,] \;=\; E[\,\theta_{x_i}\,[\,1-C(\theta_{x_{-i}})\,\gamma\,]$$

Use this expression to derive the properties of each technique summarised below:

$$P(S,S) \;=\; \theta_S\,[\,1-(1-\theta_S)\gamma\,] \;=\; 1/2 - 1/4\gamma$$

$$P(R,S) \;=\; \mu_R\,[\,1-(1-\theta_S)\gamma\,] \;=\; 1/2-1/4\gamma - (1-2\alpha_1)\Delta\,(1-\gamma/2) \;=\; P(S,S) - (1-2\alpha_1)\Delta\,(1-\gamma/2)$$

$$P(S,R) \;=\; \theta_S\,[\,1-(1-\mu_R)\gamma\,] \;=\; 1/2-1/4\gamma - (1-2\alpha_1)\Delta\,\gamma/2 \;=\; P(S,S) - (1-2\alpha_1)\Delta\,\gamma/2$$

$$P(R,R) \;=\; E[\,\theta_R\,[\,1-(1-\theta_R)\gamma\,] \;=\; 1/2-1/4\gamma - (1-2\alpha_1)\Delta + \Delta^2\,\gamma \;=\; P(S,S) - (1-2\alpha_1)\Delta + \Delta^2\,\gamma$$

Proof of lemma 3. This lemma requires comparing second period's payoffs that follow two different cases: when there is experimentation in the first period (i.e., x_1 is either R,S, S,R, or R,R); and where there is no experimentation (i.e., x_1 is S,S). It is useful to define some notation: $V_2(\mathbf{x}_1) \;\equiv\; E[\,P(x_2)\,|\,\mathbf{x}_1\,]$

i) Comparing second period's benefits under S,S vs R,S ($\emptyset$ vs h):

$$V_2(R,S) - V_2(S,S) \;\;=\;\; \mu_R\,(\theta_R^+ - \theta_S) \;-\; \mu_R\,\gamma\,(\,\theta_R^+(1-\theta_R^+) - \theta_S(1-\theta_S)\,)$$

Where I use that $prob(h=1) = \mu_R$. This expression can be rewritten as:

$$\Delta(\Delta + \alpha_1 - 1/2) \;+\; \Delta^2\,\gamma\,\mu_R \;\;=\;\; -(1-2\alpha_1)\Delta\;(1/2 + \Delta^2\gamma) \;+\; \Delta^2\,(1+\gamma/2)$$

Which is positive when

$$\alpha_1 \;>\; \frac{1}{2} \;-\; \frac{\Delta\,(1+\gamma/2)}{1+2\,\Delta^2\,\gamma}$$

ii) Comparing second period's benefit under S, R vs R, R (h_R vs h_R, h_R):

$V_2(R,R) - V_2(S,R) \;=\; prob(h_R = 1, h_R = 1)\,(\,\theta_R^{++} - \theta_S\,) \;-\; prob(h_R = 1)\,(\,\theta_R^{+} - \theta_S\,) \;+$
$(\,prob(h_R = 1, h_R = 1) - prob(h_R = 1)\,)\;\Delta^2\,\gamma$

Which is equivalent to

$$(1/4 - \Delta^2)\,[\,\Delta\,(1-2\alpha) \;-\; \Delta^2\,\gamma\,]$$

Which is positive when:

$$\alpha_1 \;<\; \frac{1}{2} \;-\; \frac{\Delta\,\gamma}{2}$$

Which is always true under assumption 1.

Proof of lemma 4. Computation just follows from making net payoffs of experimenting positive.

a) Overall net payoffs of movement from S, S to R, S are

$$P(R,S) - P(S,S) \;+\; \delta\,[\,V_2(R,S) - V_2(S,S)\,]$$

Which is simple to rewrite as:

$$-(1-2\alpha_1)\,\Delta\;\left[\,(1-\gamma/2) \;+\; (\,\Delta^2\gamma + 1/2)\,\delta\,\right] \;+\; \Delta^2(1+\gamma/2)\,\delta$$

Which is positive, meaning that there is an equilibrium with experimentation, whenever:

$$\alpha_1 \;\geq\; 1/2 \;-\; \frac{\delta\,\Delta\,(1+\gamma/2)}{2\,(\,1 \;-\; \gamma/2 \;+\; \delta/2\,(1+\,2\,\gamma\,\Delta^2)\,)} \;\equiv\; \alpha^*$$

b) Overall net payoffs of movement from S, R to R, R are

$$P(R,R) - P(S,R) \;+\; \delta\,[\,V_2(R,S) - V_2(S,S)\,]$$

Which is simple to rewrite as:

$$-(1-2\alpha_1)\,\Delta\,(1-\gamma/2-(1/4-\Delta^2)\,\delta)\;+\;\Delta^2\,\gamma\,(\,1-(1/4-\Delta^2)\,\delta\,)$$

Which is positive when:

$$\alpha_1\;\geq\;1/2\;-\;\frac{\Delta\,\gamma\,(1-(1/4-\Delta^2)\,\delta)}{2\,(\,1-(1/4-\Delta^2)\,\delta\;-\;\gamma/2\,)}\;\equiv\;\alpha_J^*\qquad\text{for }\delta\;<\;\tfrac{1-\gamma/2}{1/4-\Delta^2}$$

And

$$\alpha_1\;\leq\;1/2\;-\;\frac{\Delta\,\gamma\,(1-(1/4-\Delta^2)\,\delta)}{2\,(\,1-(1/4-\Delta^2)\,\delta\;-\;\gamma/2\,)}\;\equiv\;\alpha_J^*\qquad\text{for }\delta\;>\;\tfrac{1-\gamma/2}{1/4-\Delta^2}$$

Where the second case is excluded by assumption 1.

c) Finding $\hat{\delta}$. By comparing thresholds, we can see that $\alpha^* \;\leq\; \alpha_J^*$ (under assumption 1) whenever:

$$\delta\;\leq\;\tilde{\delta}\;\equiv\;\frac{B+\sqrt{-4(4-\gamma^2)(1/4-\Delta^2)\Delta\,\gamma\,(2-\gamma)+B^2}}{2(2+\gamma)(1/4-\Delta^2)}$$

Such that:

$$B\;\equiv\;(2-\gamma)(1+\gamma/2)-\Delta\gamma(1+2\Delta\,\gamma)$$

Proof of lemma 5.

a) CP experiments with one agent (in the first period) when:

$$P(R,S)+P(S,R)\;-\;2\,P(S,S)+2\delta\,prob(h=1)\,[P(R,R)-P(S,S\,|\,h=1\,)]\;\;\geq\;\;0$$

Which can be rewritten as:

$$-(\,1-2\alpha_1\,)\Delta\;+\;2\,\delta\;\;\big[\,-(1-2\alpha_1)\Delta\;(1/2+\Delta^2\gamma)\;+\;\Delta^2\,(1+\gamma/2)\,\big]\;\;\geq\;\;0$$

Implying that there is experimentation when:

$$\alpha_1 \;\geq\; 1/2 \;-\; \frac{\delta\,\Delta\,(1+\gamma/2)}{1\;+\;\delta\,(1+\,2\,\gamma\,\Delta^2)} \;\equiv\; \alpha^{CP}$$

b) CP experiments with two agents (in the first period) when:[23]

$$2P(R,R){-}P(R,S){-}P(S,R){+}2\delta(\,prob(h=1,h=1)\,[P(R,R){-}P(S,S\,|\,h=1,h=1\,)]\,-\\ prob(h=1)\,[P(R,R)-P(S,S\,|\,h=1\,)]\,)\;\;\geq\;\;0$$

The equation above can be rewritten as:

$$-(1-2\alpha_1)\,\Delta+2\Delta^2\,\gamma\;+\;2\delta\;\left[\,(1/4-\Delta^2)\,\Delta\,(1-2\alpha_1)\;-\;\Delta^2\,\gamma\,(1/4-\Delta^2)\,\right]$$

Implying that the CP experiments with two agents (instead of one) whenever:

$$\alpha_1\;\;\geq\;\;\frac{1}{2}\;\;-\;\;\frac{\Delta\,\gamma\,(\,1-(1/4-\Delta^2)\,\delta\,)}{1-2\,\delta\,(1/4-\Delta^2)}\;\;\equiv\;\;\alpha_J^{CP}$$

c) Using the same reasoning used in the decentralised equilibrium (lemma 4), the CP is going to experiment with two agents at the experimentation frontier whenever: $\delta < \hat{\delta}^{CP}$.

Proof of Proposition 1.

Direct computation of two cases:

i) When $\delta < min\{\hat{\delta}, \hat{\delta}^{CP}\}$ or $\delta > max\{\hat{\delta}, \hat{\delta}^{CP}\}$, the decentralised equilibrium and CP choice coincide in terms of the number of agents experimenting and comparison of thresholds shows that CP experiments for a wider range of α_1 whenever contagion is a secondary effect: $\gamma < 1$.

ii) When $min\{\hat{\delta}, \hat{\delta}^{CP}\} < \delta < max\{\hat{\delta}, \hat{\delta}^{CP}\}$, comparison is made between thresholds with different number of agents experimenting. Similar results hold.

[23]Note that whenever $\alpha_1 > 1/2-\Delta$, the first period cost for CP is smaller than for an individual agent, meaning that a single agent does not take into account the full benefits of coordination (and negative contagion externality is small).

1.7.2 Appendix of section 1.4

Proof of Lemma 5, 6 and 7

These lemmas are immediate using the properties of the expected status. Conditional on techniques $\mathbf{x} = (x_i, x_{-i})$ expected status $L(x_i, x_{-i})$ is given by:

$$E[\theta_{x_i}\,(1-\theta_{x_{-i}})(1-\gamma)]\;g \;-\; E[\theta_{x_{-i}}\,(1-\theta_{x_i})(1-\gamma)].$$

Implying the following payoffs:

$$L(S,S) \;=\; 1/4\,(g-1)\,(1-\gamma)$$

$$L(R,R) \;=\; 1/4\,(g-1)\,(1-\gamma) \;+\; \Delta^2\,(1-g)\,(1-\gamma) \;=\; L(S,S) \;+\; \Delta^2\,(1-g)\,(1-\gamma)$$

$$L(R,S) \;=\; 1/4\,(g{-}1)\,(1{-}\gamma) - (1{-}2\alpha_1)\Delta\,1/2\,(g{+}1)\,(1{-}\gamma) \;=\; L(S,S) - (1{-}2\alpha_1)\Delta\,1/2\,(g{+}1)\,(1{-}\gamma)$$

$$L(S,R) \;=\; 1/4\,(g{-}1)\,(1{-}\gamma) + (1{-}2\alpha_1)\Delta\,1/2\,(g{+}1)\,(1{-}\gamma) \;=\; L(S,S) + (1{-}2\alpha_1)\Delta\,1/2\,(g{+}1)\,(1{-}\gamma)$$

Note that $L(R,R) < L(S,R)$ when:

$$\Delta^2\,(1-g)\,(1-\gamma) \;<\; (1-2\alpha_1)\Delta\,1/2\,(g+1)\,(1-\gamma)$$

Which holds under assumption 1'.

Proof of Lemma 8

$$-(1{-}2\alpha_1)\Delta\,1/2\,(g{+}1)\,(1{-}\gamma) \;+\; \Delta^2\,(1{-}g)\,(1{-}\gamma) \;-\; \delta\;(1/4{-}\Delta^2)\;\Delta^2\,(1{-}\gamma)\,(1{-}g)$$

Implying that this net payoffs are positive when:

$$\alpha_1 \;\geq\; \frac{1}{2} \;-\; \frac{\Delta\,(1-g)\,[1-\delta\,(1/4-\Delta^2)\,]}{g+1}$$

Note that this threshold is:

- Above 1/2 whenever $\delta \; > \; \frac{1}{1/4-\Delta^2}$

- Always above $1/2 - \Delta\frac{1-g}{1+g}$. (implying that if assumption 1 and $g > 1-\gamma$ hold, this condition is never satisfied).

Proof of Lemma 9 The second period's net **status benefits** of experimentation with one agent (pioneering) are higher than the same benefits from **intrinsic utility** when:

$$\mu_R \left[\, L(R,R) - L(S,S) \,\right] \;\geq\; \mu_R \left[\, E[\, P(R,R) \,|\, h=1 \,] - P(S,S) \,\right]$$

Which is equivalent to:

$$-(1-2\alpha_1)\Delta\,[1+2\Delta^2\gamma]+\Delta^2(2+\gamma) \;\geq\; -(1-2\alpha_1)\Delta\,[\Delta^2\,(1-\gamma)(1-g)] + 1/2\Delta^2\,(1-\gamma)(1-g)$$

Which is true when:

$$\alpha_1 \;\; \leq \;\; \frac{1}{2} - \Delta\,\frac{(2+\gamma) - 1/2\;(1-\gamma)(1-g)}{2\;(\;1+\Delta^2\,(2\gamma-(1-\gamma)(1-g))\;)}$$

Or for the special case where there is no contagion ($\gamma = 0$) and $g = 0$ is:

$$\alpha_1 \;\; \leq \;\; \frac{1}{2} - \frac{3/2\;\Delta}{2\;(1-\Delta^2)} \;\; \equiv \;\; \hat{\alpha}(\Delta)$$

Where $\hat{\alpha}(0) = 1/2$ and $\hat{\alpha}(1/2) = 1/2 - \Delta$.

Proof of Proposition 2

a) Thresholds for CP. As shown above, CP will experiment with (at least one agent) when:

$$\alpha_1 \;\; \geq \;\; 1/2 \;-\; \frac{\delta\,\Delta\;(1+\gamma/2)}{1\;+\;\delta\,(1+\;2\,\gamma\,\Delta^2)} \;\; \equiv \;\; \alpha^{CP}$$

CP experiments with two agents (instead of one) when:

$$\alpha_1 \;\; \geq \;\; \frac{1}{2} \;\;-\;\; \frac{\Delta\,\gamma\;(\,1-(1/4-\Delta^2)\,\delta\,)}{1-\delta\,2\,(1/4-\Delta^2)} \;\; \equiv \;\; \alpha_J^{CP}$$

b) Threshold for status driven consumers (when ζ is large enough for decisions to be driven by status payoffs).

There is experimentation with status payoffs whenever:

$$\alpha_1 \;\; \geq \;\; \frac{1}{2} \;\; - \;\; \frac{1/2\, \delta\, \Delta}{\frac{1+g}{1-g} \; + \delta \; 2\, \Delta} \;\; \equiv \;\; \alpha^*(L)$$

c) Comparing thresholds, under δ large enough that equilibrium experimentation is done by a pioneer[24].

It can be shown that $\alpha^{CP} \;\; \geq \;\; \alpha^*(L)$, whenever:

$$g \;\; \leq \;\; \frac{\delta(1 - 4\Delta^2) - (1 + \gamma)}{\delta(1 - 4\Delta^2) + (2 + \gamma)}$$

Which is positive when $\delta \geq \frac{1+\gamma}{1-4\Delta^2}$

Or written in terms of the discount factor:

$$\delta \;\; \geq \;\; \frac{1 + \gamma \; + \; 3g \; + \; g\,\gamma}{(1 - 4\Delta^2)\; (1 - g)}$$

[24]Or, in other words, experimentation frontier is at a "pioneer equilibrium", as defined in the main text.

Chapter 2

A new explanation for pricing of subscription contracts: preventing learning of behavioural consumers.

2.1 Introduction

Subscription contracts, characterised by a fixed access fee and an agreed usage pricing, can be found in a wide range of service industries: fitness (e.g., gym memberships), communications (e.g., mobile calls and data plans), utilities (e.g., electricity contracts), banking (e.g., credit cards) and online entertainment (e.g., Netflix). Three-part tariffs is a particular structure for usage pricing in these subscription contracts where initial usage faces zero usage pricing (marginal pricing equal to zero) and later usage is highly charged (positive marginal pricing). This pricing can be found, for instance, in most mobile phone contracts where consumers pay a fixed access fee (monthly contract fee) which gives them access to usage which is priced in the following way: it is free for the first x communication/data units and it is highly charged (with overage charges) when usage crosses a given threshold (named "allowance"). Despite their prevalence, these pricing structures cannot be explained by standard contract theory models.

This paper proposes a new mechanism to explain why firms in service sectors offer subscription contracts with three-part tariff (3PT) pricing. This mechanism relies on the fact that, when choosing optimal contracts, firms are likely to be aware that consumers' behaviour differs from classical assumptions and, instead, follows three basic principles identified by economic empirical literature (and reviewed below).

First, consumers make behavioural mistakes in their contract choices by mispredicting future usage[1][2]. Second, consumers can reduce future mistakes by using their experience in previous contracts to learn. Third, consumer learning is more likely to happen when mistakes are highly visible, for instance when they are clearly reflected on their bill.

These assumptions augment previous behavioural industrial economics' literature, which analysed the firms' optimal pricing when facing behavioural consumers who make mistakes, by introducing the realistic and empirically backed assumption that consumers can use their past experience in markets to learn and avoid future mistakes. By improving our understanding of market interaction, this previously overlooked factor can allow us to improve regulatory policy in these markets. In the case developed here, consumers' learning is going to reduce the firms' profits. Hence, when the firm faces same consumers repeatedly and cannot commit to future contracts, it is going to design contracts with pricing that reduces learning opportunities for consumers which, under some conditions, leads to optimally offering subscription contracts with 3PT structure.

Previous literature. Standard contract theory models predict that the optimal marginal pricing (i.e., usage pricing) in subscription contracts is equal to the marginal cost. This result follows from the fact that, under standard conditions, subscription contracts allow the firm to first price discriminate consumers through the fixed access fee paid at the time of contracting. Using this fee, the firm can extract fully the expected consumer surplus in this contract, implying that the marginal pricing (pricing per usage) should optimally be used to induce consumers to consume at a point where this surplus is maximised - that is, it should be set to the marginal cost level.[3]

Despite this result, in real subscription contracts with 3PT, the marginal pricing is actually below marginal cost for initial units used (initial usage is subsidised)

[1]This assumption is common to both my paper and Grubb (2009).

[2]What empirical studies in several markets show is that economists, using the same ex ante (past usage) data available to consumers, would have made better choices over the available subscription contracts than consumers did, as evaluated ex post.

[3]When firms face consumers with non-contractible information about their valuation for the good, they will be required to screen consumers. This fact changes the optimal pricing but still does not generate 3PT as: (1) In standard screening, the firm still chooses allocations that maximise surplus for large types but distorts downwards the allocations of consumers with lower types. This is, however, the opposite of what 3PT do, as the distortion from screening implies higher marginal pricing for low units. (2) In sequential screening, as modelled by Courty and Li (2001), 3PT pricing can be explained only when types SOSD each other, which is contradicted by empirical observations about preferences in these markets.

and above marginal cost for later units (later usage is overcharged). Whilst there are many explanations for firms' wish to subsidise usage (such as meter effects, loss aversion, network effects, etc. for instance), these explanations cannot account for firms' decisions to simultaneously overcharge high usage. A new explanation based on consumer mistakes in forecasting future usage has been suggested by Grubb (2009). However, recent empirical evidence in the same markets (surveyed below) appears to question the validity of one the main assumptions of this model: consumers make systemic mistakes and never learn from these.

The model developed in this paper complements Grubb (2009) by studying the incentives of the firm when behavioural consumers who make mistakes can also learn from them. Under said assumption, this model is able to explain both (i) firms' motivation to offer 3PT pricing, and (ii) observed comparative statics and trends in service sectors.

Three part tariffs. 3PT are a pricing structure of subscription contracts which contains (1) a fixed fee paid independently of usage, (2) a marginal price equal to zero for the first usage units (e.g., free minutes or included data in the cellular phone market) and, (3) a positive marginal price after a threshold of usage is crossed (e.g., overage charges). Although cellular phone contracts are typical examples of usage for 3PT, the same can also be observed in contracts for services such as broadband, credit cards and car rental.

The model. The basic framework builds upon two literatures: classical contracting models and empirically motivated behavioural consumers. The standard set-up of my model consists of the following: each period, a monopolist offers a menu of one period contracts to a continuum of consumers. Consumers have a common per-period utility which is determined by their individual tastes realisations (drawn every period). Each consumer has a fixed type which can be either high or low, with equal probability. These individual types are going to affect their distribution over tastes: high type consumers are more likely to draw large tastes than low type consumers are. Each period, consumers are offered one period menu of subscription contracts (technically, sequential contracts[4]) before observing their taste realisation in that period. Given this timing, consumers choose a contract (subscription) before deciding on their their level of consumption in this subscription period which is driven by their individual taste draw. In the context of the cellular phone market example above, this implies that consumers choose their contract before learning

[4]This Sequential Contract's timing was first introduced by Courty and Li (2001).

about their call opportunities in each billing period.

The novelty of this paper's approach to this problem lies in introducing new assumptions that bring more realism to an important dimension of decision making in subscription-based contracts: how well consumers predict their future preferences.[5] Firstly, the traditional assumption about consumers' full knowledge of their own type is relaxed. Instead, I assume that, the consumers are uninformed of their type at birth. In particular, they are born with an unbiased prior about being high or low type. Secondly, I assume that, while consumers can use their memory to rationally learn their type over time, consumers' memory is imperfect. For simplicity, I assume that consumers remember only their past bills and forget their actual past consumption levels and tastes. These two key assumptions generate new dynamic incentives for firms' pricing. To capture these effects in a stationary environment, I further assume that the monopolist faces a stable population of consumers, born every period, who live for two periods only. This modelling device improves tractability and guarantees that consumers never learn their type fully over their two period lifespan.[6]

Finally, the firm's solution (the maximisation of the sum of per period discounted profits) becomes truly dynamic under two further assumptions on contracts matching the environment in many regulated service industries: (1) firms cannot use information from past contracts to build exclusive contracts[7], and (2) contracts are negotiated in every billing period (i.e, there is limited commitment). These arguably realistic assumption have important assumptions for the analysis of this problem. Any problems with future (consumer) learning could be assumed away if, instead, the firm could contract today over future periods, or if the firm could perfectly discriminate consumers after they learned. Similarly, these assumptions rule out ratchet effects (Laffont and Tirole, 1986) since the firm cannot discriminate consumers over information observed and the aggregate distribution of a continuum of consumers is known.

The main pricing trade-off. The main mechanism of this model relies on the

[5] As reviewed below, these two assumptions can also be found in Ali (2011) where behavioural consumers who face uncertainty about their own preferences and can learn about them in the context of self-control problems. In this paper, Nageeb Ali focuses mainly on characterising the consumers' decisions over commitment contracts in this context.

[6] Possible justifications for this assumption consist of rapid technological change in these industries (e.g., landlines, cellular phone services, internet services, etc.) and knowledge depreciation.

[7] This assumption, which rules out the possibility of special deals exclusive to some consumers, is without loss of generality as long as the cost of targeting consumers are high or regulation forbids this practise, both are common issues in this industry.

fact that the firms' profits are decreasing in the degree of consumer's sophistication, a concept shared with the literature of rational firms contracting with behavioural consumers (Spiegler, 2011). When the firm meets consumers repeatedly and consumers can use past data to learn about their mistakes, firms will have an incentive to try to decrease consumer learning. In particular, when consumers have an imperfect memory such that they remember predominantly their past bills, a rational firm will realise the importance of prices for consumer learning and, as a result, distort pricing away from what previous models would have predicted. This dynamic trade-off can be thought as balancing present costs of preventing learning and future costs from consumer learning.

The cost of consumer learning arises from the fact that the firm's profits are decreasing in the consumers' sophistication, which increases as consumers learn. Remember that, even though the consumers' population is heterogeneous (composed of an equal share of high and low types), the consumers are not informed about their type when they are born. The intuition for why consumer learning decreases profits is easier to understand by studying the following two extreme cases. First, consider the case where consumers are all informed about their own types. In this case, the firm - because it cannot price discriminate between types - faces an adverse selection problem and needs to screen this information by paying information rents. Second, instead, consider the case where all consumers were uninformed about their types such that there is no asymmetric information problem. In this case, the profits are higher compared to the informed case since the consumers have the same average willingness to pay but now the firm does not need to screen. More generally, when learning occurs, consumers are going to observe different information leading to a divergence in their beliefs about their likelihood of being high type and, hence, in their willingness to pay for a contract. This effect exacerbates the adverse selection problem since the firm cannot use the information learned by the consumer (it is non-contractible) to offer consumers with heterogeneous beliefs different contracts. Therefore, profits decrease with consumer learning since the firm pays higher screening costs but faces a continuum of consumers with the same average willingness to pay.

The cost of preventing learning arises from the fact that the firm needs to use flat marginal pricing (zero marginal pricing) to prevent consumer learning. When consumers use prices to learn about past events, flat prices provide the least information to consumers thus reducing consumer learning and increasing future profits. However, preventing consumer learning is going to reduce current profits, as the firm maximises current profits by offering strictly increasing allocations which can only be implemented with strictly increasing prices. When offering flat pricing, the firm distorts prices and the allocations offered, reducing the surplus it can extract from the consumers and, hence, its profits.

Results. In light of these facts, the main results of this paper characterise the optimal pricing of the firm. I show that, under some conditions, 3PT are the optimal solution to the trade-off discussed above and lead to the highest profits in this dynamic problem. As discussed above, the cost from consumer learning is the loss in future profits from screening consumers with different beliefs (about being high type). As standard in screening problems, this cost depends heavily on the consumer with the lowest willingness to pay which, in this case, are the consumers who have the lowest (posterior) belief. This suggests that the optimal way to prevent consumer learning should be to use pricing to influence this same lowest future consumer belief. I show that this can be done by manipulating pricing for the lowest quantities as observed in 3PT. In this framework, consumers with high types are more likely to consume more. As a consequence, the lower the consumption level recalled by the consumer (through bills) the lower its future belief. Therefore, the optimal way to prevent consumer learning requires distorting the lowest quantities in contracts which justifies 3PT pricing.

The optimal size of a 3PT (the length of the flat region of pricing[8]) can be analysed by thinking about the firm's decision of increasing marginally the size of the 3PT. When the firm increases the size of 3PT, it increases the cost from distorting prices (cost of preventing learning) and it decreases the information rents paid as a result of learning from consumers' prices (cost of learning). Importantly, under standard assumptions on functional forms, the profit losses from information rents are going to increase with the consumer valuation of the contracts, while the profit losses from price distortion are relatively constant. An implication of this fact is that the decision of which 3PT size to choose is going to change with the value of the good to the consumer. Therefore, the firm optimally chooses to offer increasing pricing environments (3PT of small size) in low consumer value environments; and longer 3PT (or flat) pricing in higher consumer value environments. This pricing follows common stylised facts in these markets.

[8]Within my model, I informally refer to the size of a 3PT as the length of the region (number of tastes) that is offered zero marginal pricing

2.2 Related literature.

Pricing of subscription contracts has been shown to deviate from marginal pricing in a variety of contexts as: fitness services (Della Vigna and Malmendier, 2006); landline services (Palacios Huerta and Miravete, 2016); cellphone services (Grubb and Osborne, 2015); online shopping (Goetler and Clay, 2010); internet services (Lambrecht and Skiera, 2006); etc. A particular deviation from what standard contract theory would predict happens when the firms choose to use 3PT pricing. As surveyed by Grubb (2009), this deviation is particularly puzzling since it means that firms facing positive marginal costs are offering a marginal pricing which (compared to their marginal costs): under-prices low usage and over-charges high usage. Standard contract models - where firms offer subscription based contracts, as developed by Courty and Li (2001) - predict that this price distortion is never optimal under the observed consumer characteristics in these markets. Grubb (2011) develops in some detail the argument that, because "subscription contracts" allow the firm to first price discriminate consumers (by extracting future expected consumer surplus through the entry fee[9]), its marginal prices (prices per minute in cellphone markets) should be used to maximise the consumers surplus which can be extracted with this fee. When the consumers are homogeneous, this means setting marginal prices equal to the marginal costs.[10] Even when consumers are heterogeneous with different willingness to pay but have rational expectations, the firm optimally discriminates by offering the best possible allocation to high usage consumers and by distorting low usage consumers, which is the opposite of 3PT pricing.[11]

Alternative behavioural explanations for this pricing have focused mainly on why firms offer pricing regions with zero marginal pricing (e.g., Lambrecht and Skiera, 2006). These partial explanations are commonly grouped under the label of "flat-rate bias". A notorious exception is Grubb (2009) which explains both distortions

[9]This argument relies on the classical assumption of quasi-linear preferences. This assumption is harmless as long as expenditure in these goods is a small proportion of the consumer's budget such that income effects can be ignored.

[10]Prices for quantities that do not affect consumer behaviour (for instance, consumers never choose those quantities) could be set at any level. Hence, this argument could be used to explain why firms use 3PT but also any other price shape for non-relevant price regions. While theoretically possible, this argument does not seem to be relevant in these markets as evidence reviewed in Grubb (2011) suggests that quantities over all the pricing regions discussed here are actually chosen by a significant proportion of consumers.

[11]This argument relies on the empirical fact that consumer types in these markets differ by having distributions over tastes/usage that first order stochastically dominate (FOSD) each other and not SOSD, as pointed out in Osborne and Grubb (2014) for the cellphone market. When this is the case, as shown in Courty and Li (2001), price distortions are similar to standard adverse selection problems.

of 3PT by assuming that consumers are overconfident on their ability to predict future usage. When facing such consumers the firm is going to offer 3PT to take advantage of consumers missperceptions. This model, however, relies on consumers who make systematic mistakes in predicting their tastes but can never learn from them. By opposition, my new explanation for the use of 3PTs relies on the dynamic incentives of the firm when facing consumers who make behavioural mistakes when choosing contracts but can learn over time from their past choices. This contrasts with the vast literature that studies the interaction of firms with bounded rational consumers (surveyed by Spiegler, 2011), where consumer learning is mostly absent. A single exception (to my knowledge) can be found in Nageeb Ali (2011), where naive consumers with self-control problems can learn about their mistakes and have access to commitment devices.

The novel behavioural assumptions in my model are based on recent empirical evidence in service markets that suggests that consumers do not have perfect self-awareness about their own preferences. This assumption is documented in various service industries such as: cellphone services (Grubb and Osborne (2015) and Ater and Landsman (2016)); landline services (Palacios Huerta and Miravete (2015)) and banking (Agarwal, Driscol, Gabaix and Laibson (2013) and Ater and Landsman (2013)). There are two observations, common to the papers cited above, that suggest an alternative way to model how consumers forecast their future tastes.

First, consumers do not seem to be fully informed about their own consumption preferences. Behavioural economics has long shown that the way we perceive our own preferences is imperfect. For instance, Kahneman (2011)'s model of human mind, composed of two systems, implies that our decisions are not always driven by a thoughtful and self-aware decision process. Moreover, Carruthers (2011) and other authors have argued that our ability to develop self-knowledge is not much different from our ability to "mind-read" other people (and, hence, it is poor)[12], except we have more data. Behavioural industrial economics literature (surveyed by Spiegler, 2011) has analysed extensively the implication of this assumption in several market contexts. A general result shared by most of this literature is that the firms' profits are decreasing in the consumers' sophistication (i.e. self-awareness

[12] "Mind-read" means understanding other people's emotions and motivations. These authors complement this analogy with experimental evidence that suggests that the average person performs very badly at "mind-reading" tasks.

about preferences).[13][14]

Second, over time, consumers can learn about their preferences but their memory of past events is imperfect. Economic modelling of the market implications of memory limitations is sparse[15] and the few exceptions can be found in models in the context of: search theory with costly endogenous memory (Dow (1991)); organizational change and inertia (Hirshleifer and Welch (2001)); and the Permanent Income Hypothesis' framework (Mullainathan (2002)). In this context, Piccione and Rubinstein (1997) develop a framework, used in my discussion section, to think about what happens when consumers are aware of their memory limitations. They show the necessity of introducing the concept of personal equilibrium between different selves of a consumer (i.e. the consumer at different points in time). Finally, experimental evidence suggests that not all memories are encoded in the same way. In particular, memories associated to actions seem to be retained better (e.g., paying a bill).[16].

[13] A particularly nice application on dynamically inconsistent preferences and the importance of self-awareness can be found in O'Donoghue and Rabin (2000)

[14] The approach followed in this paper differs from the one in the motivated beliefs' literature (see Benabou, 2016). In this literature, consumers' beliefs are incorrect not just because of lack of information but also because consumers actively try to influence their own beliefs for other reasons (e.g. self-image, correct for motivational issues).

[15] Although it is possible to interpret many of the behavioural biases as a consequence of limited memory, for instance availability bias.

[16] For a survey see Tulving and Thomson (1973)

2.3 Set-up

In each period $t \in \{1, 2, ...\}$, a monopolist offers a menu of *single period contracts* to a continuum of consumers. The monopolist maximises the discounted sum of per period profits Π (given by revenues P minus convex costs $C(q)$). Each single period contract (in period t) describes payments P for the provision of a non-durable good $q_t \in [0, \infty)$. The consumer population is stable and has an overlapping generation (OLG) structure: every period a new cohort of consumers of measure one is born and lives for two periods. In period t, young and old consumers are the consumers who were "born", respectively, in period t and $t-1$.

Consumers (their standard side). Tastes $\theta \in \Theta \subset \mathbb{R}$ determine the consumers' utility function. When consuming q units of the good acquired at cost p, a consumer with taste θ has utility $\quad \theta \; v(q) \;\; - \;\; p \quad$ where standard assumptions hold: $v(0) = 0$, $v'(q) > 0$ and $v''(q) \leq 0$ for all q. Tastes belong to the discrete[17] and equidistant set $\Theta \; \equiv \; \{\theta_1, \theta_2, ..., \theta_N\}$ with at least three tastes[18][19] $(N \geq 3)$ and $\theta_i - \theta_{i-1} \equiv \Delta > 0$ for all $i > 2$.

Each consumer is born with a type μ, fixed over time, which can be *high* (H) or *low* (L) with the same probability. The consumer's type $\mu \in \{L, H\}$ determines the full support distribution $f_\mu(\theta)$ from which his tastes θ are individually drawn, every period. High type consumers are more likely to draw large values of tastes, in the sense that the high type's distribution First Order Stochastically Dominates (FOSD) the low type's[20]. For tractability, we focus on the case where the distributions of the types are linear:

$$f_H(\theta_i) \;=\; \tfrac{1}{N} \;+\; \alpha \,(\, i - \tfrac{N+1}{2} \,) \qquad \text{and} \qquad f_L(\theta_i) \;=\; \tfrac{1}{N} \;-\; \alpha \,(\, i - \tfrac{N+1}{2} \,)$$

This functional form (relaxed in the discussion section) has the advantage that statistical inference is simple: high type's distribution "monotone likelihood dominates" the distribution of low types (i.e., $\frac{f_H(\theta_i)}{f_L(\theta_i)}$ is increasing in tastes θ_i)[21]; and type heterogeneity can be captured by a single (positive) parameter α. Call β the

[17] I choose discrete tastes space for ability to perform numerical simulations (many applications of dynamic mechanism design use numerical methods which require a discrete state space) and tractability of the solution. By using a discrete taste space, we can focus on cases where the monotonicity constraints are not binding. With continuous tastes instead, solutions require optimal control techniques which make profit comparisons over different contracts intractable.

[18] Less than 3 tastes would not be enough to characterise all pricing characteristics.

[19] Alternative formulation is $\theta_i \;\equiv\; \theta_1 + \Delta\,(i-1)$

[20] Such that cumulative distributions $F_\mu(\theta)$ satisfy: $1 - F_H(\theta) \geq 1 - F_L(\theta)$, for all $\theta \in \Theta$

[21] In other words, tastes work as types' signals with a likelihood ratio that is increasing and bounded.

probability that a consumer is of type H, which will be discussed in more detail below.

Contracts' structure. Since we are modelling subscription-based services (e.g., cellphone contracts), it is assumed that consumers accept their contract before observing their individual and time specific taste θ (i.e., ex ante).[22] Without loss of generality, focus on direct mechanisms where the firm offers contracts conditional on direct reports of consumer's information.[23] At the time of choosing their contract, the consumers only differ in the information they have about their individual type, characterised by β. Then, "contract β" - a contract designed for consumers with beliefs β - has $q_\beta(\theta)$ and $P_\beta(\theta)$ as the respective quantity and transfers pair designed for a consumer with taste θ. For each consumer facing such contracts, the timing of each period t can be described in three stages:

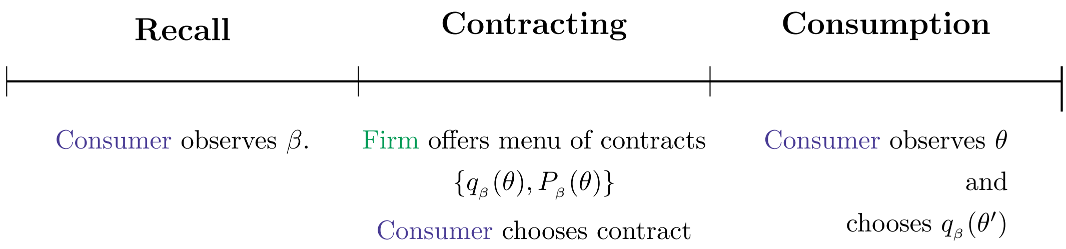

In the *"contracting stage"*, as common in this literature, the firm cannot "third price discriminate" consumers and, hence, it cannot offer contracts exclusively available to consumers with particular characteristics or history.[24]. Then, given the menu of contracts available, each consumer chooses a contract or rejects them all, getting an outside option normalised to zero. Finally, in the *"consumption stage"*, each consumer's taste θ is drawn and privately observed; and, conditional on acceptance of contract β, the consumer chooses quantity $q_\beta(\theta')$ and pays $P_\beta(\theta')$.

Consumers (their non-standard side). The following assumptions represent

[22] This represents a variation of the classical mechanism design timing where, instead, the firm offers contracts *before* the consumers observe their current period's tastes. This timing, which introduces uncertainty about future tastes at the time of contracting, is commonly used to model subscription based services (e.g., Courty and Li (2001), Miravete and Palacios Huerta (2015)).

[23] As argued by Courty and Li (2001), restrict attention to contracts that use sequential direct mechanisms, which depend on sequential reports of consumers' private information (beliefs $\beta \in \vec{\beta}$ and tastes θ), formalised as reporting strategies at contracting (CT) and consumption (CS) stage: $\sigma_{CT} : \vec{\beta} \rightarrow \vec{\beta}$ and $\sigma_{CS} : \vec{\beta} \times \Theta \times \vec{\beta} \rightarrow \Theta$. As it will be clear later, it is without loss of generality to assume that the firm cannot/does not send message to consumers with information observed, as the firm's profits are generally decreasing in consumer learning.

[24] In practise, this means that this firm would not call these consumers with special offers, either because it would be too costly or the firm was forbiden by regulatory or other reasons.

the novelty of my framework. Firstly, *young consumers are uninformed about their own individual type μ, but can learn it.* When consumers are born they have homogeneous (unbiased) priors about being of high type given by $\beta_0 = 1/2$ and, hence, attribute probability $f_{\beta_0}(\theta) \equiv (1-\beta_0)\, f_L(\theta) + \beta_0\, f_H(\theta)$ to each taste θ.[25] Based on their experience when young, consumers can learn. This experience will influence their rational expectations about their type and, hence, their expectations about their taste draw when old: an old consumer, with updated belief β of being high type, expects taste θ with probability $f_\beta(\theta) \equiv (1-\beta)\, f_L(\theta) + \beta\, f_H(\theta)$. Learning will affect how the old consumer's evaluate contracts, as the likelihood of drawing high tastes is increasing in the consumer's belief β, such that $1 - F_\beta(\theta) > 1 - F_{\beta'}(\theta)$ whenever $\beta > \beta'$.

Secondly, *consumers' memory is imperfect.* Old consumers only remember part of their history: they forget their (previously realised) taste θ, but remember their past transfers to the firm (i.e., their bill) $P_\beta(\theta)$ and the contract they were offered $P_\beta : \Theta \to \mathbb{R}$. Imperfect memory will make learning generally slower. Counterfactually, consumers with a perfect memory would use their individual taste realisations as a signal to update their probability of being high type. Instead, consumers with imperfect memory can only use their previous bills to try to rationally infer their past tastes, meaning that they have access to a noisier signal of their type. The lifetime timing of a consumer born in period t is given by:

YOUNG			**OLD**		
Prior β_0	Chooses contract β from menu of contracts	Observes privately θ and picks $q_\beta(\theta')$ $\downarrow$. Pays $P_\beta(\theta')$	Observes past memory and updates beliefs: β'	Chooses contract β'' from menu of contracts	Observes privately θ'' and picks $q''_\beta(\theta''')$ $\downarrow$. Pays $P_{\beta''}(\theta''')$

Finally, the consumers' choice in the consumption stage is assumed to be myopic. That is, consumption choices within a contract are optimal with respect to that period only. This rules out the possibility that consumers choose non-optimal quantities within a contract in an effort to communicate with their future (forgetful) self. Under some conditions, discussed further in discussion section, relaxing this

[25] In a standard model, consumers would be informed about their type, such that individual β would be either one or zero.

assumption would not qualitatively change the results.

Functional forms and welfare. I assume the functional forms introduced by Mussa and Rosen (1978) for both tractability and exposition purposes. This means that utility is linear $v(q) = q$ and costs are quadratic $C(q) \;=\; \frac{1}{2}\, q^2$, implying that ex post welfare - defined by the sum of consumer surplus (with taste θ) and firm's profits - is given by:

$$W(q, \theta) \;=\; \theta\, v(q) - C(q) \;=\; \theta\, q - \tfrac{1}{2}\, q^2$$

And the efficient quantity $q^*(\theta) = \theta$, which maximises ex post welfare, is unique, such that efficient welfare is $W^*(\theta) \equiv max_q\, W(q, \theta) \;=\; W(q^*(\theta), \theta) \;=\; \theta^2$.

2.4 Formalising the problem and preliminary steps

The main argument developed in this section is that, even though the firm offers single period contracts, its problem is actually going to be dynamic. I develop this point by showing that consumer learning with an imperfect memory is going to introduce a link between the profits achievable in any period t and the pricing offered in period $t-1$. This is a simple consequence of two facts shown in this section: beliefs in period t are a function of past $(t-1)$ pricing choices; and the profits in any period depend on the current vector of consumers' beliefs.

2.4.1 Beliefs and profits: formalising the firm's problem.

Let $\vec{\beta}_t = \{\beta_1, \beta_2, ..., \beta_J\}$, such that $\beta_j > \beta_{j-1}$, be the ordered set of beliefs in period t[26] (of young and old consumers), with density and cumulative probabilities functions given by $g(\beta)$ and $G(\beta)$, respectively.[27] When facing a consumer belief vector $\vec{\beta}$, the firm offers (single period) menu of contracts $\{P_\beta(\theta), q_\beta(\theta)\}$ that solve the following constrained optimisation problem:

Problem 1.

$$Max_{q_\beta(\theta)\geq 0\,,\,P_\beta(\theta)} \quad \sum_{\beta\in\vec{\beta}_t} E\ \left[\ P_\beta(\theta_i) - C(q_\beta(\theta))\mid \beta\ \right]\ \ g(\beta)\ +\sum_{s=1}^{\infty}\ \delta^s\ \ E\left[\ \Pi(\ \vec{\beta}_{t+s}\,)\mid P_t\ \right]$$

subject to: *(for all posteriors $\beta\ ,\ \beta'\ \in \vec{\beta}_t$)*

$$\begin{array}{lll} \textbf{\textit{Ex post}}\ \ \boldsymbol{IC_{\theta\,|\,\beta}} & U_\beta(\theta,\theta)\ \geq\ U_\beta(\theta,\theta') & ,\ \forall\theta,\theta'\in\Theta \\ \textbf{\textit{Ex ante}}\ \ \boldsymbol{IC_\beta} & E[\,U_\beta(\theta)\mid\beta\,]\ \geq\ E[\,U_{\beta'}(\theta)\mid\beta\,] & \\ \textbf{\textit{Ex ante}}\ \ \boldsymbol{PC_\beta} & E[\,U_\beta(\theta)\mid\beta]\ \geq\ 0 & \end{array}$$

Where the last term $(\sum_{s=1}^{\infty}\ \delta^s\ \ E\left[\ \Pi(\ \vec{\beta}_{t+s}\,)\mid P_t\ \right])$ represents the firm's continuation value. The constraints of this problem reflect that, as usual, the direct mechanism has to be incentive compatible and individually rational. Let $U_\beta(\theta,\theta')$ be the surplus of such a consumer (in a contract β) with taste θ' who reports a

[26] Subscript t will be dropped whenever it is clear

[27] Given our assumptions on the probability distribution of the types, the belief set $\vec{\beta}$ is enough to identify the distribution function $g(\beta)$. Hence, frequently notation I will refer to the relationship between profits and the belief vector without mentioning the implied probability distribution.

taste θ at a later stage, such that $U_\beta(\theta',\theta) \equiv \theta' v(q_\beta(\theta)) - P_\beta(\theta)$. Importantly, while the *incentive compatibility (IC)* constraints are similar to standard screening models, now they need to hold both ex post and ex ante.[28] Since the firm cannot offer exclusive contracts to consumers with different beliefs, the firm needs to make consumers self-select ex ante: $E[U_\beta(\theta) \mid \beta] \geq E[U_{\beta'}(\theta) \mid \beta]$ for any $\beta, \beta' \in \vec{\beta}_t$. Then, consumers who accept contract β ex ante are still going observe their taste privately and, hence, the firm needs to provide incentives for consumers with different tastes to self-select: $U_\beta(\theta) \equiv U_\beta(\theta,\theta) \geq U_\beta(\theta,\theta')$ for all $\theta, \theta' \in \Theta$.[29] The consumer's participation decision is made ex ante, meaning the *participation constraint (ex ante PC)* only needs to hold in expectation: $E[U_\beta(\theta) \mid \beta] \geq 0$ for all posteriors $\beta \in \vec{\beta}_t$.

Finally, it is useful to define some terminology on other characteristics of the solution, which are related to other implicit constraints of this problem. We say that no shut-down constraints (also called non-negativity) are satisfied when all quantity solutions of this problem, arrived at while ignoring this constraint, are non-negative. We say monotonicity in tastes or beliefs is satisfied when quantities offered in optimal contracts, again while ignoring these constraints, are respectively weakly increasing in tastes or beliefs.

2.4.2 Learning: linking pricing today and beliefs tomorrow

I discuss here the relationship between old consumers' beliefs in period $t+1$ and the prices charged to young consumers $P_{\beta_0}(\theta)$ in period t. Pricing is not going to change the average beliefs that the firm faces, since rational learning (martingale property) implies that the average belief of old consumers $E[\beta] = \sum_j \beta_j\, g(\beta_j)$ is equal to the prior β_0. However, pricing will affect the distortion of consumer beliefs. To formalise this fact, define some notation on the old consumer's posteriors:

Definition 1. *Let $\beta(\theta_i)$ be the posteriors of a consumer who inferred, from prices, that his past taste realisation was θ_i. Similarly, let $\beta(\Theta')$ be the posterior when consumer could only infer that his past taste belongs to sub-set Θ'.*

This definition allows us to discuss more formally what happens to future beliefs when a firm changes its price today. Imperfect memory is create a relationship between pricing and beliefs. To understand this relationship, consider first the

[28] Refer to decisions ex ante and ex post, as decisions before and after the consumers observes his (time and individual specific) taste θ.

[29] This ex post IC of consumer β will not depend on consumer's beliefs β since beliefs do not enter utility function and, hence, do not matter for ex post constraints.

simpler counterfactual where consumers' memory is perfect (i.e., they can remember their past tastes), such that old consumers' have heterogeneous beliefs given by $\beta(\theta_i)$ for each observed taste θ_i that satisfy:

$$\beta(\theta_1) < \beta(\theta_2) < ... < \beta(\theta_N)$$

As a result of larger tastes being more likely to be drawn by high type consumers, by definition. Then, note that consumers with imperfect memory only remember their past transfer to the firm P and the contract they were offered $P_{\beta_0} : \Theta \to \mathbb{R}$. The consumers can try to infer their previous taste using the inverse function: $P_{\beta_0}^{-1} : \mathbb{R} \to \Theta$ whenever it exists.[30] If pricing is strictly increasing in all tastes (i.e., $P_{\beta_0}(\theta_{i+1}) > P_{\beta_0}(\theta_i)$ for all $i = 1, 2, ..., N-1$), consumers can do this operation and, hence, they get the same information from prices as they would get from observing past tastes. However, this is not always the case. For example, when the firm offers fully "flat" pricing for all tastes (such that $P_{\beta_0}(\theta_{i+1}) = P_{\beta_0}(\theta_i)$ for all $i = 1, 2, ..., N-1$), the consumers are not going to learn anything about their types. This means that their posterior is $\beta(\Theta) = \beta_0$. To discuss this issue, use the following definition of pricing:

Definition 2. *A characterisation of the pricing offered to young consumers (with belief β_0) can be written in terms of tastes: $\bar{\Theta}(P_{\beta_0}) = \{\Theta_k\}_{k=1}^K$ is a partition of taste set Θ where in each block Θ_k the prices are constant/flat: $P_{\beta_0}(\theta_i) = P_{\beta_0}(\theta_{i-1})$ for all $\theta_i, \theta_{i-1} \in \Theta_k$.*

For example, when $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$ and $\bar{\Theta}(P_{\beta_0}) = \{\{\theta_1, \theta_2, \theta_3\}, \{\theta_4\}\}$ this implies that pricing is $P_{\beta_0}(\theta_1) = P_{\beta_0}(\theta_2) = P_{\beta_0}(\theta_3) < P_{\beta_0}(\theta_4)$.

Then, use this definition to say that a pricing P'_{β_0} is a flattening of pricing P_{β_0} when $\bar{\Theta}(P_{\beta_0})$ is a refinement of $\bar{\Theta}(P'_{\beta_0})$. This allows us to show what happens to $t+1$ beliefs when the firm flattens the pricing of sub-sets of tastes in period t (such that sub-sets that used to be offered strictly increasing pricing are now offered flat pricing):

Proposition 1. *Because of imperfect memory, the dispersion of beliefs in period t is a function of pricing in $t-1$. In particular, flattening pricing decreases dispersion of future beliefs and, hence, decreases consumer learning.*

Proof. When the firm offers flat pricing for a sub-set of tastes Θ', it aggregates the signals the consumers would have observed if they could remember each taste:

[30] In general, inverse function might not exist and it might be the correspondence (when remembering P and $P_{\beta_0} : \Theta \to \mathbb{R}$) $\bar{\Theta}(P) = \{\theta \in \Theta : P_{\beta_0}(\theta) = p\}$

$\beta(\Theta') = E[\beta(\theta) \mid \theta \in \Theta']$. This fact implies that the distribution of beliefs in the second period originated by flattening pricing of a contract second order stochastically dominates the original distribution of beliefs. □

Finally, a corollary of this proposition is that the firm can use current period's pricing to influence the lowest belief of old consumers in the next period. This property will be important since the consumers with lowest belief are not paid information rents and, hence, play an important role on the firm's profits. To show this result, use the fact that consumers who get the lowest taste draws are, by rational learning[31], going to have lowest future beliefs and focus on pricing in this region:

Definition 3. *For any (one period) menu of contracts A, define $\#(A)$ as the counting variable which yields the number of tastes in the lowest flat region of the pricing of offered to young consumers (contract β_0): $\#(A) = argmax_{\{i\}}\, P_{\beta_0}(\theta_i)$ for $P_{\beta_0}(\theta_i) = P_{\beta_0}(\theta_1)$.*[32]

This means that, when contract A does not start with a flat area, the counting variable takes the value one ($\#(A) = 1$) and when contract A starts with a flat region of size l, the counting variable takes the value l ($\#(A) = l$). This notation can be used to show that the firm can use first period's prices to influence the lowest posterior:

Lemma 10. *The lowest belief in $t+1$ generated by (period t) contract A increases with $\#(A)$.*

Proof. The lowest belief comes from the consumers with the lowest taste realisations in the first period, and for contract A is given by:

$\beta_1 = E[\beta(\theta) \mid \theta \in \{\theta_1, ...\theta_{\#(A)}\}]$ which is increasing in $\#(A)$. □

[31] As the likelihood ratio of types is increasing in tastes.

[32] $\#(A) = \{\, \theta_{\#(A)} \,:\, P_{\beta_0)}(\theta_1) = P_{\beta_0}(\theta_2) = ... = P_{\beta_0}(\theta_{\#(A)}) \,<\, P_{\beta_0}(\theta_{\#(A)+1}) \,\}$

2.5 The basic trade-off

Solving the dynamic problem of the firm, described in the previous section, requires understanding the trade-off between present and future profits. The dynamic trade-off arises from the fact that, while consumer learning in period t decreases period $t+1$'s profits, preventing this same consumer learning is costly, as it decreases period t's profits. This section considers *separately* the two factors that drive this trade-off: adverse selection and price distortions. As it is clear below, adverse selection is the mechanism that links consumer learning and the firm's profits, and price distortions link the efforts of preventing consumer learning and profits. By evaluating the impact on profits of each of these factors, we can develop a deeper understanding of the mechanisms at stake, before approaching the full problem.

2.5.1 The cost of learning (adverse selection)

Define the cost of learning as the loss in (single period) profits from facing informed consumers (that is, consumers with heterogeneous belief vector $\vec{\beta}$); instead of uninformed consumers, with the initial prior β_0. Formally, this cost can be defined as:

Definition 4. *Cost of learning (C_L) is a measure of the loss in profits that comes from consumer learning:*

$$C_L(\vec{\beta}) \;\equiv\; \hat{\Pi}(\beta_0) \;-\; \hat{\Pi}(\vec{\beta})$$

Where $\hat{\Pi}(\vec{\beta})$ are the highest profits that can be achieved in a given period when facing belief vector $\vec{\beta}$, or (formally) the solution to problem 1 when $\delta = 0$ and consumer beliefs are $\vec{\beta}$.

Solving for profits. To compute the profits that compose the cost of learning, note that these are the solution to a static mechanism design problem (i.e., the static *Problem 1*). Below, using standard contract theory literature steps, I solve for the profits with consumer learning:

First, show that the firm has to pay information rents. The firm will have to provide positive surplus to consumers in order to make them self-select into contracts (i.e., information rents), since this is a typical adverse selection problem: consumers have non-contractible information (their beliefs), which influences their individual valuation of contracts, since:

Lemma 11. *Consumers' expected utility from any contract is increasing in their belief β: $E[U_{\beta'}(\theta)\,|\,\beta\,]$ is increasing in β.*

Proof. A standard manipulation of the expected utility equation gives us:

$$E[U_{\beta'}(\theta)\,|\,\beta\,] \;=\; U_{\beta'}(\theta_1) + \textstyle\sum_{i=2}^{N} \left(U_{\beta'}(\theta_i) - U_{\beta'}(\theta_{i-1}) \right) (1 - F_\beta(\theta_{i-1}))$$

Which is increasing in β given the following two facts:

i) Consumer surpluses $U_\beta(\theta)$ are increasing in tastes θ, since ex post IC binding implies that surpluses' difference is given by $U_\beta(\theta_i) - U_\beta(\theta_{i-1}) \;=\; \Delta\, v(q_\beta(\theta_{i-1})) \;>\; 0$ (for any contract $\beta \in \vec{\beta}$ and for all $\theta_i \in \Theta$)[33].

ii) $F_\beta(\theta_{i-1})$ is decreasing in β, since consumers with larger beliefs - i.e., who are more likely to be high types - are also more likely to draw large tastes, by FOSD property of types. □

Second, show how information rents influence the firm's profits. Note that the profits accrued from each contract β_j can be written as:[34]

$$E[\,P_\beta(\theta) - C(q_\beta(\theta))\,|\,\beta_j\,] \;=\; E[\,W(q_\beta(\theta))\,|\,\beta_j\,] \;-\; IR(\beta_j)$$

That is, profits depend only on two factors: the welfare generated by this transaction; and the information rents (IR) paid to the consumers, where information rents (IR) for consumers with belief β_j are $IR(\beta_j) \;=\; E[\,U_{\beta_j}(\theta)\,|\,\beta_j\,]$. Finally, to characterise information rents, use the fact that (when monotonicity constraint defined above holds) only the local downward IC will hold, such that the information rents will depend on local deviations:

Lemma 12. *When quantities $q_\beta(\theta)$ are increasing in β, the firm pays information rents to screen consumers with different beliefs for all consumers with beliefs above β_1:*

$$IR(\beta_j) \;-\; IR(\beta_{j-1}) \;=\; \sum_{i=2}^{N} \Delta\, v(q_{\beta_{j-1}}(\theta_{i-1}))\,(F_{\beta_{j-1}}(\theta_{i-1}) - F_{\beta_j}(\theta_{i-1})) \quad \textit{for } j > 1$$

[33] Given that consumers' utility satisfies single crossing, consumers with larger tastes have higher (marginal and absolute) utility from the good and, hence, need to be offered higher surpluses and quantities $q_\beta(\theta)$ are increasing in θ_i. This is a standard step (first suggested by Mirrlees (1971)), which means that the "global" ex post IC constraints hold if and only if monotonicity in θ and "local" downward ex post IC hold.

[34] Using a simple manipulation of profit function.

and $\quad IR(\beta_1) \; = \; 0$

Which also means that consumers with the lowest belief are not paid any information rents (i.e., $IR(\beta_1) = 0$). Putting these steps together, it is easy to show that:

Proposition 2. *The highest profits achievable in any period when the firm faces consumers with beliefs $\vec{\beta}$ are given by:*

$$\hat{\Pi}(\vec{\beta}) \;=\; \textstyle\sum_{\beta_j \in \vec{\beta}} E[\, W^*(\, \hat{\theta}_i(\beta_j)\,) \,|\, \beta_j \,] \;\; g(\beta_j)$$

Whenever monotonicity in θ and β and no shut-down constraints do not hold, the normalised tastes are:

$$\hat{\theta}_i(\beta_j) \;\equiv\; \theta_i - X_i(\beta_j) \;\;=\;\; \theta_i \;-\; \tfrac{F_{\beta_j}(\theta_i) - F_{\beta_{j+1}}(\theta_i)}{f_{\beta_j}(\theta_i)} \, \tfrac{1-G(\beta_j)}{g(\beta_j)} \, \Delta$$

A consequence of the proposition above is that, when there is no consumer learning (i.e., all the old consumers are still uninformed), profits are simply given by:

$$\hat{\Pi}(\beta_0) \;=\; E[\, W^*(\theta) \,|\, \beta_0 \,]$$

Which is a direct consequence of the fact that there is no adverse selection problem and the firm can extract all the consumer surplus by offering a take-it-or-leave-it contract.

Cost of learning. The proposition above shows that profits when facing consumers with heterogeneous beliefs are equivalent to profits when the firm faces homogeneous consumers with lower taste levels[35] . In this sense, define $X_i(\beta_j)$ as the normalised information rents (for belief β_j and taste θ_i), since this term reflects the loss in profits generated by information rents, when these loses are integrated into the profit function. Hence, the cost of learning[36] depends crucially on the value of normalised information rents:

$$C_L(\vec{\beta}) \;=\; \hat{\Pi}(\beta_0) - \hat{\Pi}(\vec{\beta}) \;\;=\;\; \textstyle\sum_j \; E\big[\, W^*(\theta_i) - W^*(\, \theta_i - X_i(\beta_j)\,) \;\, | \; \beta_j \,\big] \; g(\beta_j)$$

This equation also makes clear that the relationship between consumers' beliefs and the cost of learning is complicated: C_L depends non-linearly on normalised information rents (since the welfare function is convex) and normalised information

[35] The fact that we can get this simple result is a consequence of linearity of utility.
[36] When monotonicity constraints are not binding.

rents have a non-linear relationship with beliefs.[37]

Properties of the C_L. The following are important facts about the cost of learning, which play a crucial role in the future analysis of optimal contracts.

PROPERTY 1. The cost of learning is a result of beliefs being non-contractible.

As the C_L's expression makes clear, it is not consumer learning per se that decreases profits, but the fact that the learned information is non-contractible. If the firm could offer exclusive contracts conditional on beliefs, then it would not need to pay information rents such that $X_i(\beta_j) = 0$ for all j and i. In this case, the profits with learning would equal the profits without learning, since, by the martingale property:[38]

$$\sum_j E[W^*(\theta_i) \mid \beta_j]\, g(\beta_j) \;=\; E[W^*(\theta_i) \,|\, \beta_0] \;=\; \hat{\Pi}(\beta_0)$$

The next property shows that, under some conditions, the firm's profits respond mainly to the proportion of consumers with the lowest belief. That is, for the cost of learning, the composition of belief vector will not matter as much as the proportion of consumers who are being paid these rents.

PROPERTY 2. The cost of learning, when the type heterogeneity is small, is mainly determined by the proportion of consumers with belief β_1, that is consumers who are not paid information rents.

This property relies on the fact that, when the type heterogeneity is small, the problem can be simplified considerably. This can be shown using the steps overviewed below:

First, the normalised information rents $X_i(\beta_j)$ can be shown to decrease when the type heterogeneity α lowers. This is a result of two complementary effects: information rents for any beliefs decrease, as the underlying distributions of the two types are more similar (i.e., $F_L(\theta_i) - F_H(\theta_i)$ decreases); and consumers learn less information from the same past observations, leading to less informed beliefs (i.e., $\beta_{j+1} - \beta_j$ decreases). Rewriting normalised information rents allows us to see this effect:

$$X_i(\beta_j) \equiv \quad \frac{F_{\beta_j}(\theta_i)-F_{\beta_{j+1}}(\theta_i)}{f_{\beta_j}(\theta_i)}\, \frac{1-G(\beta_j)}{g(\beta_j)}\, \Delta \;=\; (\, \beta_{j+1}-\beta_j \,)\, \frac{F_L(\theta_i)-F_H(\theta_i)}{f_{\beta_j}(\theta_i)}\, \frac{1-G(\beta_j)}{g(\beta_j)}\, \Delta$$

[37] Note that the equation above can be rewritten as $C_L(\vec{\beta}) \;=\; E\left[\int_{\theta_i - X_i(\beta_j)}^{\theta_i} \frac{\partial W^*(\theta)}{\partial \theta}\, \mathrm{d}\theta \mid \beta_j\right]$ using a simple application of martingale property and fundamental theorem of calculus.

[38] Hence, this is a consequence of two facts of profits of a monopolist without adverse selection: profits are given by welfare, which is linear in beliefs; and rational learning implies that the expected value of the posteriors is equal to the prior: $\sum_j \beta_j\, g(\beta_j) \;=\; \beta_0$ (Martingale property).

Second, the cost of learning depends on information rents in a non-linear fashion. However, this effect of information rents on the cost of learning can be well approximated by a linear function when normalised information rents are small. Focusing on the expected value of the linear effect of information rents, we can see that it is not going to depend on the composition of beliefs. Instead, it is only going to depend on the prior belief β_0 and the lowest belief β_1 (and so indirectly on the proportion of consumers paid information rents):

$$\sum_{\beta_j} E[\, X_i(\beta_j) \,|\, \beta_j \,]\; g(\beta_j) \;=\; (\, \beta_0 \,-\, \beta_1 \,)\;\Delta \sum_i (F_L(\theta_i) - F_H(\theta_i))$$

Where this relationship follows from the fact that beliefs are bounded by rational learning, such that their expected value is equal to their initial prior. This argument leads to the following lemma:

Lemma 13. *When type heterogeneity is small enough such that $\alpha < \bar{\alpha}$: the normalised information rents $X_i(\beta_j)$ are small, which implies that the influence of consumer beliefs in the cost of learning comes only through the lowest belief β_1 (where $\beta_1 = min\{\beta_j\}$). That is, when $\alpha < \bar{\alpha}$ it can be shown that:*

$$C_L(\vec{\beta}) \;\approx\; (\, \beta_0 \,-\, \beta_1 \,)\;\Delta \sum_i (F_L(\theta_i) - F_H(\theta_i))\;\; 2\,\theta_i$$

An immediate consequence of this lemma is that the cost of learning will increase in the level of tastes (i.e., the value of θ_1) that the firm faces:

PROPERTY 3. The cost of learning is increasing in the level of tastes.

The intuition for this property follows from the reasoning behind information rents. The firm pays information rents because (at the time of contracting) it faces consumers with heterogeneous private beliefs that affect their willingness to pay for contracts. As shown before, the willingness to pay for a contract (expected utility) depends on beliefs for two reasons: the probability of getting large taste draws increases with the probability that a consumer is high type (i.e., β); and, in any (implementable) contract, the consumers always get (ex post) surpluses which are increasing in their taste realisation. For instance, in contract β, the consumer's surplus gain from getting a taste realisation θ_i instead of θ_{i-1} is always given by:

$$U_\beta(\theta_i) - U_\beta(\theta_{i-1}) \;=\; \Delta\; v(q_{\beta}(\theta_i))$$

Finally, it is simple to see that, as quantities offered are increasing in the taste level, the difference between surpluses of consumers with high beliefs and low beliefs increase with the overall taste level, as measured by θ_1. This implies that the information rents increase in the taste level.

2.5.2 The cost of preventing learning (price distortions)

Define the cost of preventing learning as the difference in profits from using pricing to prevent consumer learning. Formally, using the intuition develop before, we say that the firm prevents consumer learning when it offers flat pricing for tastes. Then, we can say that the firm prevents consumer learning of multiple sub-sets $\{\Theta_k\}_{k=1}^{K}$ when it sets pricing such, for each k, pricing is flat such that $P_{\beta_0}(\theta) = P_{\beta_0}(\theta')$ for all $\theta, \theta' \in \Theta_k$. To isolate this effect, ignore the adverse selection problem by considering the case where the firm faces a single consumer with known belief, such that $\vec{\beta} = \beta$. In particular, since the firm only wants to prevent learning of young consumers, focus on the case where $\vec{\beta} = \beta_0$. Then, the cost of preventing learning is defined as:

Definition 5. *The cost of preventing learning with pricing $\bar{\Theta}(P_{\beta_0})$ is defined as the difference between the profits when not preventing learning and the profits when preventing learning:*

$$C_P(\bar{\Theta}(p)) \;\; = \;\; \hat{\Pi}(\,\beta_0\,) - \;\; \hat{\Pi}(\,\beta_0\,|\,\bar{\Theta}(p)\,)$$

Where $\hat{\Pi}(\,\beta\,|\,\bar{\Theta}(P_{\beta_0})\,)$ are the highest single period profits, when the firm faces consumers with a single belief β_0 and prevents learning of subsets using pricing defined by $\bar{\Theta}(P_{\beta_0})$. Formally, this is the solution to problem 1 with $\delta = 0$ and pricing constraints given by:

$$P_{\beta_0}(\theta) = P_{\beta_0}(\theta') \qquad \forall \theta\ ,\ \theta'\ \in \Theta^P(k) \quad \forall k$$

Computing profits. This problem can still be solved using standard mechanism design steps, as follows. First, note that to satisfy the new pricing constraint (for contract β_0), the firm has to pool quantities for all tastes in each subset Θ_k, such that $q_{\beta_0}(\theta) = q_{\beta_0}(\theta')$ for $\theta, \theta' \in \Theta_K$. This happens because consumers with different tastes, who are charged the same prices, only report their tastes truthfully (IC) if they are offered the same quantities.[39]

[39]Importantly, this new constraint is always going to be binding. The reason for this is that, when this constraint is absent, the optimal solution for this problem relies on prices increasing

Second, the firm does not need to pay information rents, given the absence of an adverse selection problem. Hence, the profits are proportional to the welfare generated by the allocations offered. Define $\mu_k \equiv E[\theta \mid \theta \in \Theta_k]$ as the average taste for each subset k prevented from learning. Then, since utility is linear in tastes such that $E[\theta\ v(q) \mid \theta \in \Theta_k] = \mu_k\ v(q)$, the problem of what quantity to offer to each taste sub-set Θ_k can be further simplified:

$$max_{\bar{q}_k}\ \ E[\, W(\bar{q}_k\, ,\, \theta) \mid \theta \in \Theta_k\,] \ \ = \ \ max_{\bar{q}_k}\ \ W(\bar{q}_k\, ,\, \mu_k)\ Prob(\Theta_k)$$

Using this equation, the optimal profits can be divided into two sub-groups: when preventing learning from a sub-set of tastes Θ_k, profits are $W^*(\mu_k)\, Prob(\Theta_k)$ for all tastes in the sub-set; when the firm is not preventing learning, the profits are given the efficient profits $W^*(\theta) f_{\beta_0}(\theta)$, for each taste:[40]

Lemma 14. *The profits when preventing learning of sub-sets $\{\Theta_k\}_{k=1}^{K}$ are given by:*

$$\hat{\Pi}(\,\beta \,|\, \Theta_k\,) \ \ = \ \ \sum_k\ Prob(\theta \in \Theta_k)\ \ W^*(\mu_k) \ \ + \ \ \sum_{\theta \notin \Theta_k}\ f_{\beta_0}(\theta)\ W^*(\theta)$$

Which implies that, as seen before, absent adverse selection, the profits when not preventing learning are given by: $\hat{\Pi}(\,\beta\,) \ = \ E[\, W^*(\theta) \,|\, \beta\,]$. This implies that the cost of preventing learning can be written in a very intuitive way, which illustrates its link to the loss in welfare from distorting prices (and, hence, allocations):

Lemma 15. *The cost of preventing learning is given by:* $\quad C_P(\Theta_k) \ \ = \ \ \sum_k\ \sigma^2(\Theta_k)\ Prob(\Theta_k)$

Where $\sigma^2(\Theta_k)$ is conditional variance given by $\sum_{\theta_i \in \Theta_k} (\,\theta_i - \mu_k\,)^2\ \frac{f_{\beta_0}(\theta_i)}{Prob(\Theta_k)}$.

The expression for the cost of preventing learning is derived from the fact that the cost of preventing learning comes from the loss on consumer's utility from distorting quantities away from the efficient level and, under the functional form assumed, the cost of allocation distortions is quadratic, since:

$$W^*(\theta) - W(q, \theta) \ = \ (\, q - q^*(\theta)\,)^2$$

and quantities are linear in tastes since $q^*(\theta) = \theta$. The equation in the lemma above implies a couple of properties which are immediate:

in tastes. Without adverse selection problem, the firm finds it optimal to implement the strictly increasing efficient allocation $q^*(\theta)$, which can only be done using strictly increasing prices.

[40] This means that the profits when preventing learning of each sub-set Θ_k are given by: $\sum_k\ Prob(\theta \in \Theta_k)\ \ W^*(\mu_k) \ \ + \ \ \sum_{\theta \notin \Theta_k}\ W^*(\theta)\ f_{\beta_0}(\theta)$

PROPERTY 1. The C_P is increasing and convex[41] *in number of tastes prevented from learning (i.e., included in subsets Θ_k)*

Which reflects the increasing costs of distorting quantities that are present whenever the problem is concave. In particular, in this case the cost of quantity distortions is quadratic.

PROPERTY 2. The C_P is scale free in terms of tastes, that is $\frac{\partial C_P(\Theta_k)}{\partial \theta_1} = 0$.

This means that when tastes change but the firm's cost of distorting quantities are constant, the C_P does not change. Importantly, when considering different levels of tastes in later analysis, the properties of the distribution functions are assumed to be constant over the taste indexes. Then, to assess how the C_P changes with taste parameters, it is useful to note that *level of tastes* can be defined by θ_1 since tastes are equidistant: $\theta_i = \Delta(i-1) + \theta_1$. Then, the properties of the firm's problem of preventing learning are similar when tastes change but the distance between tastes does not.[42]

[41] In the meaning that the function is convex, when approximated as a continuous function. Alternatively, the rate of change of C_P (as more tastes are included in flat regions) grows with the number of tastes in flat regions.

[42] This property can be generalised in a weaker version as long as the effect of tastes on C_P are bounded, that is whenever $\frac{\partial^3 W^*(\theta)}{\partial \theta^3}$ is bounded.

2.6 Optimal contract in two period environment

Assume that there is only one generation, meaning that this becomes a simple two period model. This section characterises the optimal contract which maximises profits for this case. The results developed below focus on the characteristics of the first period's pricing of the optimal contract. This is the relevant pricing since, in this stylised two period model, only the first period's pricing decision reflects the new trade-off between maximising current profits and preventing future consumer learning.

Optimal contracts and 3PT. Define optimal contracts as the short term contracts that maximise the firm's profits over two periods. Solving for the overall profit level over the two periods can be done by taking advantage of the previous solutions which isolated the effect of adverse selection and price distortion:

Lemma 16. *Profits in period t (when monotonicity constraints hold), given belief vector $\vec{\beta}_t$ and pricing choice $\bar{\Theta}_t$, are given by:*

$$\Pi_1 + \Pi_2 \;\equiv\; (\,\hat{\Pi}(\beta_0) - C_P(\bar{\Theta}_t)\,) \;+\; (\,\hat{\Pi}(\beta_0) - C_L(\vec{\beta})\,)$$

Where $\vec{\beta}$ is a function of pricing in the first period. (delete?)

Where the equilibrium profits are simplified by the fact that in the second (and last) period the firm is never going to prevent consumer learning (as implicit in the formula above). Similarly, since all consumers are born in the same period, in the first period, they share the same belief β_0 and there is no adverse selection problem. Hence, in this case, the adverse selection problem and the price distortions happen in different periods, so we don't need to worry about their interaction. We can find the optimal contract by comparing profits generated by different first period's pricing, assuming that the firm makes an optimal choice in the second period. At the optimal first period's pricing the total profits are maximised, or (as defined above) the sum of first period C_P and second period C_L are minimised, as the lemma above illustrates.

Define the 3PT of size J as a contract with pricing given by: flat pricing for the first J tastes ($\#\,(3PT(J)) \;=\; J$); and strictly increasing pricing for all the tastes afterwards:[43]

[43] Note that we allow $1 \leq J \leq N$.

Definition 6. *A contract β has pricing with 3PT (of size J) structure when $P_\beta(\theta_i) = P_\beta(\theta_{i-1})$ for all $i \leq J$ and $P_\beta(\theta_i) > P_\beta(\theta_{i-1})$ for all $i > J$.*

Then, we can show that, when comparing menus of contracts, any menu where the young consumer's contract has pricing different from 3PT is dominated by a menu where, instead, the pricing in young consumers contract has a 3PT structure:

Proposition 3. *When type heterogeneity is small enough (in particular, $\alpha < \bar{\alpha}$), contracts for young consumers have always 3PT pricing.*

This result relies on two facts shown previously. First, preventing learning is costly. That is, the cost of preventing learning is always increasing in the areas of pricing that are prevented from learning[44], as shown in property 1 of cost of preventing learning. This means that the firm is going to choose carefully which prices it distorts. Second, the firm benefits mainly from preventing learning of consumers who are going to have lowest future beliefs. This is a consequence of the fact that the (second period's) cost of learning depends mainly on the lowest belief whenever type heterogeneity is small enough (property 2 of cost of learning). Preventing learning of such consumers can only be done by offering flat pricing for the lowest quantities in a contract, as consumers use their own consumption to learn about their types and higher types consume higher quantities. When distorting pricing is this way, the firms are implicitly offering 3PT to young consumers.

Characterising 3PT. The optimal size of 3PTs can also be described by using the findings of previous sections. Consider the impact on profits of offering (to young consumers) a 3PT with a marginally larger size. That is, in our case, the impact on profits of going from a 3PT of size J to size $J+1$. Using our previous decomposition of profits into efficient welfare and costs, we can see that increasing the size of 3PT generates two opposing effects, it leads to:

- A larger number of prices distorted, which increases the cost of preventing learning: call this effect $\Delta\, C_P(J) \ < 0$.
- A lower proportion of consumers being paid information rents, which decreases the cost of learning: call this effect $\Delta\, C_L(J) \ > 0$.

[44] In other words, the cost of learning is always increasing in the number of tastes included in the sub-sets of pricing that are flat

The optimal size of 3PT are, therefore, going to depend on the properties of the cost of learning and cost of preventing learning, as the size of 3PT change. Importantly these properties are going to depend on the taste level in this market, as shown in the following lemma.

Lemma 17. *The valuation of consumers for the good (i.e., their taste level) is going to impact how costs change when increasing marginally the size of 3PT, as follows:*

i) $\Delta C_L(J)$ *is going to increase in the level of tastes in* θ_1

ii) $\Delta C_P(J)$ *is independent of the level of tastes* θ_1

While the loses from price distortions do not scale up for larger tastes, the loses from adverse selection problem grow as the level of tastes increases. This fact is more general than the present functional forms suggest.

The cost of learning comes from screening of consumers with different beliefs ex ante. As the level of tastes increase the level of information rents that the firm needs to pay are going to increase, as the general level of payoffs are higher.

The cost of preventing learning comes from the loss in profits from pooling quantities by distorting prices. This loss is not proportional to the level of tastes because by assumption tastes are equidistant, meaning that when tastes scale up their distance is kept constant. This means that this cost is not going to increase when the properties of the welfare generated by this contract does not change much.p[45]

This fact means that consumers' tastes are going to change the firm's trade-off when minimising the sum of the cost of learning and the cost of preventing learning, implying the following optimal contracts:

Proposition 4. *In optimal contract, the pricing for young consumers is:*

i) 3PT of size 1 (fully increasing) *when* $\theta_1 < \underline{\theta}_1$

ii) 3PT of size J (for $1 < J < N$*)* *when* $\underline{\theta}_1 \leq \theta_1 \leq \underline{\theta}_1$

iii) 3PT of size N (fully flat) *when* $\bar{\theta}_1 < \theta_1$

Where we can use properties of the cost of learning and preventing learning to show that this problem is concave (costs are convex). The following graphs illustrate possible costs of learning and costs of preventing learning[46] for different sizes of 3PT. In figure **??**, you can see the case where taste level is intermediate, such that the firm

[45]Technically, this can be shown to be related to whether the third derivative of the welfare function is large or not.

[46]An imaginary line joins what are the discrete points at different sizes of the 3PT. This line helps to make the results more visual.

minimizes the total costs by offering a 3PT with an interior size (that is, $1 < J < N$. For lower levels of tastes the costs of learning are smaller and the graph instead looks like figure 2. While the cost of distorting prices is constant over tastes, the cost of adverse selection increases (both in absolute and marginal value) with the level of tastes.

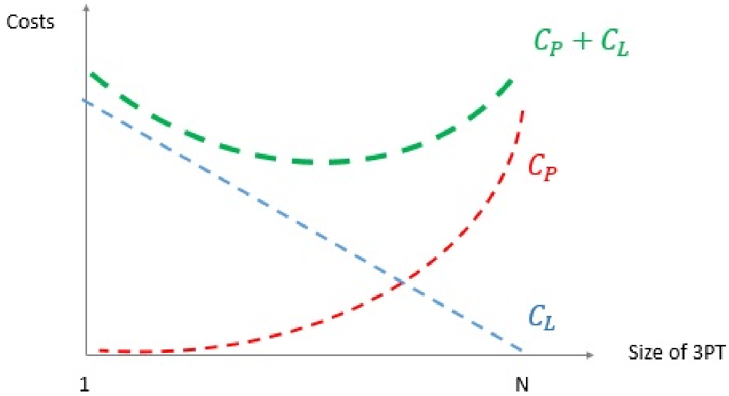

Figure 2.1: The costs of learning and preventing learning with intermediate taste levels

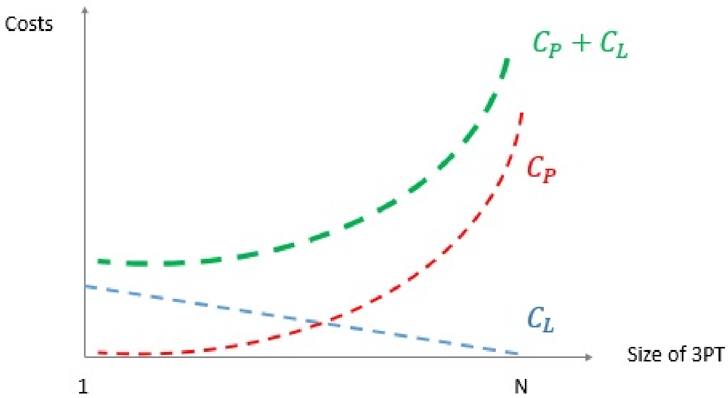

Figure 2.2: The costs of learning and preventing learning with low taste levels

Discussion This section suggests a relationship between the value of the good for the consumers and the pricing offered that seems to match observed cellphone contracts which are: increasing for low value contracts; and 3PT (with size $1 < J < N$) or fully flat for high value contracts.

Finally, note that there are two market observations that this model does not seem to be able to capture: first, the firms in reality seem to offer a limited number of contracts, instead of a contract for each belief. Second, the firm offers 3PT only in a single contract (to the young consumers).[47] In the next chapter, I extend the current the model to address these and other issues that are important in the service industry and show that, under some conditions, the present results are going to be robust.

[47]Note that the firm does not have an incentive to prevent learning of old consumers with beliefs $\beta \neq \beta_0$ and hence, their pricing is strictly increasing.

2.7 Appendix

2.7.1 Appendix of section 2.4

Meaning that, in any contract A, the lowest posterior is increasing in $\#(A)$.

Proof of prop.1 First note that the firm can only flatten pricing of convex sets of tastes, since only weakly increasing pricing (in tastes) is implementable in this framework. Then, by flattening the pricing in a contract designed for young consumers (consumers with beliefs β_0) for a sub-set Θ'', the firm changes beliefs that arise from tastes in that subset from $\beta(\theta_i)$ for each θ_i in this sub-set to $\beta(\Theta'') = E[\beta(\theta_i)\,|\,\theta_i \in \Theta''\,]$.

Proof of lemma 10. Beliefs depend positively on the tastes the consumer remember, as the likelihood ratio $L_i \equiv \frac{f_H(\theta_i)}{f_L(\theta_i)}$ is increasing in θ_i, by definition, and the posterior $\beta(\theta_i) = \frac{L_i}{1+L_i}$, derived using bayes rule, is increasing in this likelihood ratio.

This implies that the lowest belief of current belief vector comes from consumers with the lowest taste realisations in the previous period: in contract A, where the firm offers flat pricing for the first $\#(A)$ tastes, the lowest belief is:

$$\beta_1 \;=\; \beta(\theta_1, ..., \theta_{\#(A)}) \;= E[\beta(\theta)\,|\,\theta \in \{\theta_1, ...\theta_{\#(A)}\}\,].$$

2.7.2 Appendix of section 2.5

Proof of Lemma 11/12 and proposition 2

Note that lemma 11 (formula for information rents) is equivalent to saying that local downward ex ante IC holds. Below we show that both lemma 11 and proposition 2 are satisfied when monotonicity constraints holds.

Define a "relaxed solution" to the optimal solution to the problem 1 with $\delta = 0$ by incorporating local downward ICs into the original profits function to write the commonly called "virtual surplus" that the firm can extracted, as defined below. Importantly, we say that monotonicity constraint in tastes and beliefs hold when quantities implied by the relaxed solution described above $q_\beta(\theta)$ are increasing in θ and β.

Using these definitions, the usual strategy to solve this problem is to start by optimising the relaxed problem and use that solution to show to check if monotonicity constraints hold. In this spirit, rewrite the problem as:

Problem 2. Optimal quantities and prices $\{q_\beta(\theta), P_\beta(\theta)\}$ are the solution to:

$$\max_{\substack{q_\beta(\theta)\geq 0 \\ q_\beta(\theta) \;\; \textit{non decreasing in } \theta \textit{ and } \beta}} \quad E\,[\;\Phi_j(q(\theta)\,,\,\theta)\;]$$

Where $\Phi_j \equiv E[\,W(q\,,\,\theta_i - X_i(\beta_j))\,|\,\beta_j\,]$ and $X_i(\beta_j) = (IR_j - IR_{j-1})\ \frac{1-G(\beta_j)}{g(\beta_j)\,f_{\beta_j}(\theta_i)}\ \frac{1}{v(\,q_{\beta_j}(\theta_{i-1})\,)}$

Use the following observations to note that if relaxed solutions satisfies monotonicity constraints then all other constraints from original problem are satisfied:

1. As shown by Mirless (1971), when consumers satisfy single crossing, the global ex post IC can be substituted by the downward local ex post IC and monotonicity constraint in tastes.

2. Lemma 2 implies that only PC of β_1 is binding and all other PC are going to hold.

3. Generally, local downward ex ante IC is sufficient for the ex ante IC, but not necessary and sufficient (as shown by Courty and Li, 2001). However, when monotonicity in beliefs holds (that is, $q_\beta(\theta_i)$ is increasing in β), the local downward ex ante IC is necessary and sufficient condition for ex ante IC to hold, as shown below.

   a) Start by noticing that expected surplus in contract β for consumer with belief β_j can be written as:

$$E[U_\beta(\theta)\,|\,\beta'] = \; U_\beta(\theta_1) \;+ \sum_{i=2}^{N} \;\big(U_\beta(\theta_i) - U_\beta(\theta_{i-1})\big)\;(1 - F_{\beta'}(\theta_{i-1}))$$

   b) Use the fact that ex post IC binds locally and properties of preferences: $U_\beta(\theta_i) - U_\beta(\theta_{i-1}) \;=\; U_\beta(\theta_i, \theta_{i-1}) - U_\beta(\theta_{i-1}) \;=\; \Delta\;v(q_\beta(\theta_{i-1}))$ to show that expected surplus from a deviation is given by:

$$E[U_\beta(\theta)\,|\,\beta'] - E[U_\beta(\theta)\,|\,\beta] = \sum_{i=2}^{N} \;\Delta\,v(\,q_\beta(\theta_i)\,)\;(F_\beta(\theta_{i-1}) - F_{\beta'}(\theta_{i-1})\,)$$

c) Using the previous steps, the global IC can be rewritten as:

$$E[U_{\beta}(\theta)\,|\,\beta] - E[U_{\beta'}(\theta)\,|\,\beta'] \;\;\geq\;\; E[U_{\beta'}(\theta)\,|\,\beta] - E[U_{\beta'}(\theta)\,|\,\beta']$$

Where assuming that local downward ICs and substituting that fact into this equation allows us to show that it holds when monotonicity constraint holds.

Finally, note that this solution implies that the highest profits achievable in the second period are: $\hat{\Pi}(\vec{\beta}) \;=\; \sum_{\beta_j} E[\,W^{*}(\hat{\theta}_i)\,|\,\beta_j\,]$, where $\hat{\theta}_i \equiv\; max\{\theta_i - X_j^i\,,\,0\}$.

Proof of lemma 13 and property 2 and 3 of C_L.

When type heterogeneity is small, the problem can be simplified considerably.

1) When type heterogeneity is small α, the normalised information rents are small. Rewrite normalised information rents as:

$$X_i(\beta_j) \equiv \quad \frac{F_{\beta_j}(\theta_i)-F_{\beta_{j+1}}(\theta_i)}{f_{\beta_j}(\theta_i)} \, \frac{1-G(\beta_j)}{g(\beta_j)} \, \Delta \quad = \quad (\, \beta_{j+1}-\beta_j \,) \, \frac{F_L(\theta_i)-F_H(\theta_i)}{f_{\beta_j}(\theta_i)} \, \frac{1-G(\beta_j)}{g(\beta_j)} \, \Delta$$

For small enough α, it is always the case that $\frac{\partial X^i_j}{\partial \alpha} = X^i_j \left(\frac{\frac{\partial (F^i_L - F^i_H)(\beta_{j+1}-\beta_j)}{\partial \alpha}}{(F^i_L - F^i_H)(\beta_{j+1}-\beta_j)} - \frac{\frac{\partial f^i_{\beta_j}}{\partial \alpha}}{f^i_{\beta_j}} \right) > 0$ since, by definition, $(F^i_L - F^i_H)(\beta_{j+1} - \beta_j) \to 0$ when $\alpha \to 0$ while $f^i_{\beta_j}$ converges to a constant $(f^i_{\beta_0})$ when α tends to zero.[48]

2) Then, using the fundamental theorem of calculus, decompose the effect of the normalised information rents X^i_j on the cost of learning into a linear component; and an extra term that corrects for convexity (when monotonicity constraints do not hold):

$$C_L(\vec{\beta}) \;=\; \textstyle\sum_j E\left[\; \frac{\partial W^*(\theta_i)}{\partial \theta} \, X_i(\beta_j) \;-\; \int_{\theta_i - X_i(\beta_j)}^{\theta_i} \frac{\partial^2 W^*(\theta)}{\partial \theta^2} \, \big(\theta - (\, \theta_i - X_i(\beta_j) \,)\big) \, \mathrm{d}\theta \;\mid\; \beta_j \; \right] \; g(\beta_j)$$

In our case, this can be simplified to:

$$C_L(\vec{\beta}) \;=\; \sum_j E\,[\;\; X_i(\beta_j) \;\; (\, 2\,\theta_i \;\; - \;\; X_i(\beta_j) \,) \;\mid\; \beta_j \;\;] \;\; g(\beta_j)$$

And use the fact that as normalised information decrease the convex term decreases faster since the function is quadratic and so this term can be ignored for sufficiently small normalised information rents.

3) Finally, to understand the remaining linear term, note that the expected value of normalised information rents does not depend on any consumer belief other the lowest one. As beliefs are bound by rational learning, the expected value of normalised information rents is given by:

$$\sum_{\beta_j} E[\, X_i(\beta_j) \,|\, \beta_j \,] \; g(\beta_j) \;\;=\;\; (\, \beta_0 \,-\, \beta_1 \,) \;\; \Delta \sum_i (F_L(\theta_i) - F_H(\theta_i))$$

[48]NOTE THAT: $\frac{\frac{\partial (F^i_L - F^i_H)(\beta_{j+1}-\beta_j)}{\partial \alpha}}{(F^i_L - F^i_H)(\beta_{j+1}-\beta_j)}$ is not zero when α tends to zero. $\frac{\frac{\partial (F^i_L - F^i_H)(\beta_{j+1}-\beta_j)}{\partial \alpha}}{(F^i_L - F^i_H)(\beta_{j+1}-\beta_j)} = \frac{\frac{\partial F^i_L - F^i_H}{\partial \alpha}}{F^i_L - F^i_H} + \frac{\frac{\partial \beta_{j+1} - \beta_j}{\partial \alpha}}{\beta_{j+1} - \beta_j}$

Proof of Lemma 14 and Property 1 and 2 of C_P.

- The formula for the cost of preventing learning can be derived using: $W(q, \theta) = W^*(\theta) - (\, q - q^*(\theta)\,)^2$ and that $q^*(\theta) = \theta$ and $\bar{q}_k = \mu_k$. This means that the cost of preventing learning can be rewritten as $\sum_{\theta_i \in \Theta_k} (\, \theta_i - \mu_k \,)^2 \; f_{\beta_0}(\theta_i)$, which is proportional to the formula of conditional variance.

- Property 1 and 2 of C_P are a direct consequence of variance formula of uniform distribution and equidistant tastes.

2.7.3 Appendix of section 2.6

Proof of Lemma 16.

Properties of cost of learning and cost of preventing learning imply that:

i) $\Delta C_L(1) \approx \Delta C_L(N) \approx \Delta C_L(J)$ but they are increasing in the level of tastes θ_1. As it has been shown that The cost of learning decreases (almost) linearly with the level of 3PT. To show this fact use that the assumptions on type's distributions imply that $\beta(\theta_j) = \frac{f_H(\theta_i)\, 1/2}{f_{\beta_0}(\theta_i)} = \frac{1}{2} - \frac{N}{4}(N+1)\alpha + i\,\alpha\frac{N}{2}$ which is linear in i: $\Delta C_L(1) \approx \Delta C_L(N)$ Finally, note that the rate at which C_L decreases depends on the level of tastes, and these terms are increasing in the taste level: $\frac{\partial\, \Delta C_L(J)}{\partial\, \theta_1} > 0$ for all J.

ii) $\Delta C_P(1) < \Delta C_P(N)$ but they are independent of the level of tastes θ_1. As it has been shown that the cost of preventing learning is increasing and convex in the size of 3PT, but does not depend on the level of tastes: $\Delta C_P(1) < \Delta C_P(N)$ and $\Delta C_P(J)$ is independent of θ_1.

Proof of Proposition 3. (Optimal contract)

0) In "extra lemma 1" and "extra lemma 2" (below), I show that the monotonicity constraints are going to hold for the contracts/pricings that are claimed to be optimal in this proposition (i.e., 3PT, flat contracts and strictly increasing contracts). As obvious, the profits decrease when these constraints bind. Hence, the fact that they might fail for other contracts/pricings (not included in this proposition) is not a problem for optimality.

1) Using properties of cost of preventing learning and costs of learning it is immediate that total costs are convex, and hence the problem is concave.

2) Optimal size of 3PT J^* follows from optimisation:

$$J^* = \begin{cases} 1 & \text{When } \Delta C_L(1) + \Delta C_P(1) > 0 \\ J & \text{When } \Delta C_L(1) + \Delta C_P(1) < 0 \text{ and } \Delta C_L(N) + \Delta C_P(N) > 0 \\ N & \text{When } \Delta C_L(N) + \Delta C_P(N) < 0 \end{cases}$$

This solution can be rewritten using that $\Delta C_L(J)$ does not depend on J as:

$$J^* = \begin{cases} 1 & \text{When } |\Delta C_L| < \Delta C_P(1) \\ J & \text{When } \Delta C_P(N) > |\Delta C_L| > \Delta C_P(1) \\ N & \text{When } |\Delta C_L| > \Delta C_P(N) \end{cases}$$

Finally, using properties of these costs, this solution is:

$$J^* = \begin{cases} 1 & \text{When } \theta_1 < \underline{\theta}_1 \\ J & \text{When } \underline{\theta}_1 < \theta_1 < \bar{\theta}_1 \\ N & \text{When } \bar{\theta}_1 < \theta_1 \end{cases}$$

Proof of Extra lemma 1.

Relaxed solution satisfies monotonicity in **beliefs** in period t when pricing of young consumers in $t-1$ (which generated beliefs $\vec{\beta}_t$) is:

i) Strictly increasing: $P_{\beta_0}(\theta_i) > P_{\beta_0}(\theta_{i-1})$ for all i.

ii) Constant over tastes / fully flat: $P_{\beta_0}(\theta_i) = P_{\beta_0}(\theta_{i-1})$ for all i

iii) $3PT(J)$: $P_{\beta_0}(\theta_i) > P_{\beta_0}(\theta_{i-1})$ for all $i \leq J$ and $P_{\beta_0}(\theta_i) = P_{\beta_0}(\theta_{i-1})$ for all $i > J$.

Relaxed solution satisfies monotonicity in β whenever:

$$X_i(\beta_{j+1}) \;\leq\; X_i(\beta_j) \;\text{ for all }\; \theta_i \;\text{ and }\; \beta_j.$$

To show when $X_i(\beta_j)$ is decreasing in beliefs, split the expression of normalised information rents terms into two groups:

$$X_i(\beta_j) \;=\; \underbrace{(\,F_L(\theta_i) - F_H(\theta_i)\,)\,\Delta}_{\text{Independent of beliefs}} \;*\; \underbrace{\frac{\beta_{j+1}-\beta_j}{f_{\beta_j}(\theta_i)}\,\frac{1-G(\beta_j)}{g(\beta_j)}}_{\text{Depends on beliefs}}$$

The last term $\frac{1-G(\beta_j)}{g(\beta_j)}$ is a function of the pricing in the previous period offered to each taste θ_i and the proportions of consumers with each taste $f_{\beta_0}(\theta_i)$. Starting with the simplest case (case i), notice that after strictly increasing prices all consumers remember their past tastes θ_i implying that beliefs are $\beta(\theta_i)$ (as defined in the text). In this case, probability distribution just reflect probabilities of each taste:

$$g(\,\beta(\theta_i)\,) \;=\; f_{\beta_0}(\theta_i)$$

Linearity of probability functions allows us to show that posteriors are simply given by: $\beta(\theta_i) \;=\; \frac{1}{2} \;-\; \frac{N}{4}\,(N+1)\,\alpha \;+ i\,\alpha\,\frac{N}{2}$ meaning that the expression above in this case can be simplified to:[49]

$$\frac{\beta(\theta_{k+1})-\beta(\theta_k)}{f_{\beta(\theta_k)}(\theta_i)}\,\frac{1-G(\,\beta(\theta_k)\,)}{g(\,\beta(\theta_k)\,)} \quad = \quad \alpha\,\frac{N}{2}\,\frac{1}{f_{\beta(\theta_k)}(\theta_i)}\,\frac{(N-k)/N}{1/N}$$

Which is decreasing in k as α decreases since $\lim_{\alpha\to 0} f_{\beta(\theta_k)}(\theta_i) \;=\; f_{\beta_0}(\theta_k) \;=\; 1/N$ as $f_{\beta_0}(\theta_i)$ is uniform.

To analyse $3PT(J)$ case, note that only one normalised information rent $X_i(\beta_1) \;=$

[49] wher e $f_{\beta_j}(\theta_i) \;=\; \frac{1}{1/2+2\alpha(k-\frac{N+1}{2})(\beta(\theta_k)-1/2)}$

$X_i(\,\beta(\theta_1,...,\theta_J)\,)$ is different from the analysis in case i). Use linearity of probability distributions to show that:

$$\beta(\theta_1,...,\theta_J) \;=\; \frac{\beta(\theta_1)+\beta(\theta_J)}{2} \;=\; \frac{1}{2} \;-\; \frac{N}{4}\,(N+1)\,\alpha \;+\; \frac{1+J}{2}\,\alpha\,\frac{N}{2}$$

Implying that

$$\frac{\beta(\theta_{J+1})-\beta(\theta_1,...,\theta_J)}{f_{\beta(\theta_k)}(\theta_i)}\,\frac{1-G(\,\beta(\theta_J)\,)}{g(\,\beta(\theta_1,...,\theta_J)\,)} \;=\; \alpha\,\frac{1+J}{2}\frac{N}{2}\,\frac{1}{f_{\beta(\theta_k)}(\theta_i)}\,\frac{(N-J)/N}{1/N}$$

Which is simple to show that is always above $X(\beta_2)$.

Finally monotonicity is trivially satisfied when there is only one belief in t following flat pricing in period $t-1$, that is case ii).

Extra lemma 2.

Relaxed solution satisfies monotonicity in tastes θ whenever: $\theta_{i+1} - X_{i+1}(\beta_j) \geq \theta_i - X_i(\beta_j)$ for all θ_i and β_j.

This is always satisfied as $\alpha \to 0$ since $X_{i+1}(\beta_j)$ tends to zero.

Proof of Proposition 4. Contracts for young consumers have always 3PT pricing.

Proof. All contracts which are not 3PT are going to be dominated by 3PT, since:

1) Under low type heterogeneity, the interaction between C_P and beliefs is small, so the cost of preventing learning has same properties as before.

2) Show this by contradiction. Suppose a contract A is optimal and it is not a 3PT. Then, we can always show that 3PT of size $\#(A)$ has larger profits, since it has:

- C_L which is arbitrarily close to one of contract A. By lemma 4, when α is small, the C_L can be approximated by a function that depends on the lowest belief only. As the two contracts have the same lowest belief, their C_L is close!

- A smaller C_P than contract A. This argument relies now on C_P going down when flat regions above disappear. There are two effects: first, the pure effect that shows up in property 1 (C_P is increasing in tastes) and the second effect which is the fact that $\hat{\theta}$ changes as beliefs change. The problem with this effect is that it can go both ways. However, this effect is decreasing in α, so for α small enough, the result will be true. □

Chapter 3

Extensions to a new explanation for pricing of subscription contracts: preventing learning of behavioural consumers.

3.1 Introduction

This chapter augments the set-up of chapter 2 with three different extensions which establish the robustness of the original model to various considerations in this market and address potential weaknesses.

First, the original two period model is extended to a fully dynamic overlapping generations model (OLG). This is a reduced form way to consider how the mechanism identified before can be generalised to a more dynamic environment while not allowing consumers to fully learn their type. This assumption can be justified by the existence of turn-over in the market or some form of knowledge depreciation.

Second, the previously unrestricted menu of contracts (offered ex ante) is assumed to be exogenouly limited to a finite and small number of contracts. This assumption moves us closer to contracts observed in service markets and can be justified with menu costs on the firm's side or complexity-avoidance preferences on the consumer's side. For tractability, I focus on two particular cases: when the firm can only offer a single ex ante contract (section 3.1); and when the firm can offer up to two contracts (section 3.2). These new constraints change the problem of the firm considerably. By losing the ability to screen all types of consumers, the firm now faces a new trade-off, akin to the textbook monopolist's problem, between the proportion of consumers it serves and the surplus it can extract from consumers. Nevertheless, under some conditions, dispersion of consumer tastes is going to decrease profits[1] creating a decreasing relationship between profits and consumer learning and implying that the firm is still going to offer contracts with 3PT structure. Additionally, this new extension addresses possible limitations of the results of the previous chapter arising from the fact that, from the multitude of contracts offered by the firm[2], only the contract offered to young consumers had a 3PT pricing structure.

Third, the previously monopolistic set-up is augmented by introducing differentiated competition. Note that, under perfect competition, firms would never offer 3PT pricing. The reason is that, when preventing consumer learning, the firm is making a costly investment: it reduces profits today by distorting current prices, in order to generate higher future profits. When competition drives future profits to zero, firms will never choose to prevent consumer learning. However, in more realistic models of competition, as developed in section 4, firms compete but have some market power. We show that as long as firms have enough market power, they will prevent learning and, under some conditions, offer 3PT.

[1] In a mechanism similar to Johnson and Myat (2006).

[2] There is one contract for each belief β.

3.2 Optimal contract in OLG case

This section discusses the overall optimal contract in the overlapping generations (OLG) model developed in the set-up of chapter 2, focusing on the properties of pricing for the contract β_0 chosen in equilibrium by young consumers. As before, I focus on this contract as this is the only contract for which the firm has an incentive to prevent consumer learning. As discussed below, this is a limitation of the model in terms of fitting the observed contracts. Section 3 is going to address these issues by introducing other realistic constraints on the contracting space.

3.2.1 Total profits

Taking advantage of the previous solutions for profits, which isolated the effect of adverse selection and price distortion, we can now solve for the overall profit level in each period:

Lemma 18. *Profits in period t, given belief vector $\vec{\beta}_t$ and pricing choice $\bar{\Theta}_t$, are given by:*[3]

$$\hat{\Pi}(\beta_0) \quad - \quad C_L(\vec{\beta}_t) - \quad \hat{C}_P(\bar{\Theta}_t) \;\; g(\beta_0)$$

This solution relies on the fact that the lowest information rents do not depend on pricing constraints. This means that the equilibrium information rents are the same as derived (in chapter 2) for the cost of learning. Hence, we can rewrite the original problem as if the firm faces lower tastes $\hat{\theta}$, which are equal to the original tastes θ subtracted by the normalised information rents X. In this transformed problem, the restriction on prices are going to imply restrictions on quantities implementable, as before. Given linearity in tastes, the cost of preventing learning with prices is going to be equal to the variance of tastes faced. However, because now the firm pays information rents, the cost of learning is equal to the variance of normalised tastes. That is, the cost of preventing learning is:

$$\hat{C}_P(\bar{\Theta}) \quad = \quad \sum_k \, \sum_{\theta_i \in \bar{\Theta}_k} (\, \hat{\theta}_i - \hat{\mu}_k \,)^2 \; f_{\beta_0}(\theta_i)$$

Which follows the same formula of cost of preventing learning from before but for normalised tastes.

[3] Although ignored $g(\beta_0)$ might depend on t as well for some examples of previous pricing.

3.2.2 Results

Optimal contracts.[4] Solving for the profits in equilibrium is simplified by the fact that: the consumers are myopic; and per period profits depend solely on a state variable - the belief vector $\vec{\beta}$ - which follows a Markov process: it is fully determined by the firm's pricing in the previous period. It is useful to focus on *stationary contracts* where the firm offers the same pricing every period and the initial the beliefs are consistent with pricing:

Definition 7. *In stationary contracts, pricing and allocation pair are constant over time (such that $P^t_{\beta}(\theta) = P_{\beta}(\theta)$ and $q^t_{\beta}(\theta) = q_{\beta}(\theta)$); and initial beliefs are given by $\vec{\beta}$ which is a consequence of $\bar{\Theta}(P_{\beta_0})$.*

This last assumption on beliefs makes sure that the solution for stationary state contract does not depend on initial ad-hoc belief conditions. The problem of choosing between stationary contracts is simpler as the firm just needs a simple expression that depends only on the cost of learning and cost of preventing learning:

Lemma 19. *The discounted sum of profits in steady state contract is simply given by $\frac{\Pi(\vec{\beta})}{1-\delta}$ where*[5]*:*

$$\Pi(\vec{\beta}) \quad = \quad \hat{\Pi}(\beta_0) \quad - \quad C_L(\vec{\beta}) - \quad \hat{C}_P(\bar{\Theta}) \;\; g(\beta_0)$$

Importantly, restricting the contract space to stationary contracts is without loss of generality under the following conditions:

Lemma 20. *When $\delta \to 1$ and α is small, the optimal stationary contract is equal to the optimal contract.*[6]

Which rely on the fact that, for small enough α, the normalised tastes are close to original tastes implying that $\hat{C}_P(\bar{\Theta})$ is close to $C_P(\bar{\Theta})$. More importantly, this means that the current cost of preventing learning is (for our purposes) independent of the current beliefs, meaning that the choice of pricing in any period t is stationay as it influence current $C_P(\bar{\Theta})$ and the discounted $\delta\ C_L(\vec{\beta})$ in period $t+1$.

Under these conditions, we can show that the problem is equivalent to the two period model developed in the last paper, such that:

[4]Define optimal contracts as the short term contracts that maximise the firm's discounted profits.
[5]Whenever the monotonicity constraints/global IC are not violated
[6]For the same initial beliefs!

Lemma 21. *For low α, the level of tastes (general valuation of the good) θ_1 is going to influence the costs of increasing marginally the size of 3PT. The relationship is such that:*

i) $\Delta C_P(J)$ *is independent of* θ_1

ii) $\Delta C_L(J)$ *is increasing in* θ_1

Using these properties of the costs, we can show that:

Proposition 5. *In optimal contract*[7]*, when type heterogeneity α is low the pricing for young consumers is:*

i) 3PT of size 1 (fully increasing) *when* $\theta_1 < \underline{\theta}_1$

ii) 3PT of size J (for $1 < J < N$*)* *when* $\underline{\theta}_1 \leq \theta_1 \leq \underline{\theta}_1$

iii) 3PT of size N (fully flat) *when* $\bar{\theta}_1 < \theta_1$

Discussion There are two stylised facts in these market that this model does not seem to be able to capture: (1) the firms only offer a limited number of contracts, instead of a contract for each possible belief; (2) typically, low value contracts offered have strictly increasing pricing, and high value contracts have 3PT (with size $1 < J < N$) or fully flat pricing. For instance, a provider of cellphone services typically offers pay-as-you-go contracts for low value phones and 3PT or "unlimited contracts" for higher value phones. In this model, the firm only faces an incentive to offer 3PT to the young consumers, which are not necessarily the highest value consumers.[8] In the next section, I modify the model to incorporate these issues and show that our previous results are still going to hold with a contract environment that is closer to the one observed in real contracts.

[7] Assuming either steady state contracts or discount rate close to one.

[8] Note that to old consumers with beliefs $\beta \neq \beta_0$ there is no incentive to prevent learning and, hence, the pricing is generally strictly increasing.

3.3 Limited contracts

Consider the case where the firm can only offer a limited number of contracts ex ante.[9] I analyse the two simplest cases: when the firm is restricted to offer a single ex ante contract (section 3.1); and when it can offer two contracts (section 3.2). This constraint is binding since, in its absence, the firm would optimally offer a different ex ante contract for each belief β. Under this new assumption, we show that the firm is still going to choose to offer contracts with 3PT structure, under similar conditions to the previous section.

3.3.1 One contract

The firm is now restricted to offering a single ex ante contract, but consumers can still, at any point in time, have heterogeneous beliefs (characterised by a belief vector $\vec{\beta}$).[10] Beliefs are, as before, positive linked to consumers willingness to pay for a contract[11].

Therefore, the firm faces a new trade-off, akin to the textbook monopolist's problem, between the proportion of consumers it serves and the surplus it can extract from consumers. As the firm is now unable to screen consumers ex ante information, the decision of how much consumer surplus to extract is going to be equivalent to deciding on a belief cutoff $\bar{\beta}$ of consumers served by that the firm. To serve all consumers, the firm needs to charge the willingness to pay of the consumers with the lowest willingness to pay (i.e., lowest belief). Alternatively, the firm can exclude some consumers to be able to extract a higher surplus from the remaining consumers. When firm serves consumers with beliefs above $\bar{\beta}$, by charging the expected surplus of consumers with this belief, profits are given by:

$$\sum_{\beta_j \geq \bar{\beta}} g(\beta_j) \;\; E[\, W(q(\theta), \theta) \,|\, \bar{\beta}\,]$$

This problem can be analysed using the same two benchmark measures (cost of preventing learning and a cost of learning), developed in the original model. Each of these measure captures the loss in profits caused, respectively, by learning (the adverse selection problem) and preventing learning (the price distortion problem).

[9] This variation can be justified by unmodelled reasons as menu costs on the firm's side or complexity-avoidance preferences on the consumer's side.

[10] As consumers can learn using the imperfect memories from the past as before.

[11] As lemma 2 of chapter 2 still holds

Cost of learning. As the notation above makes clear, in this restricted contract environment, the profits (and, hence, this cost) are going to depend on which consumers the firm serves. When consumers are uninformed, it is optimal to offer a single contract, such that there is no need to define threshold belief served $\bar{\beta}$. [12] When consumers are informed, the highest per period profits when serving beliefs equal or above the threshold $\bar{\beta}$ are given by $\hat{\Pi}(\vec{\beta}, \bar{\beta})$. Hence, the cost of learning (given by the difference between profits in two benchmark cases: when consumers are uninformed and when they are informed) can be written as:

$$C_L(\,\vec{\beta}\,,\,\bar{\beta}\,) \;\equiv\; \hat{\Pi}(\beta_0) - \hat{\Pi}(\vec{\beta}, \bar{\beta})$$

Since the firm does not pay information rents, allocations $q(\theta)$ are going to be chosen to maximise the welfare generated by the contract. Maximising the profits found using equation **??**, implies that the cost of learning is:

$$C_L(\,\vec{\beta}\,,\,\bar{\beta}\,) \;=\; \sum_{\beta_j<\bar{\beta}} g(\beta_j)\,)\; E[\,W^*(\theta)\,|\,\beta_0\,] + \sum_{\beta_j\geq\bar{\beta}} g(\beta_j)\,)\;(\beta_0 - \bar{\beta})\;\frac{\partial\,E[\,W^*(\theta)\,|\,\beta\,]}{\partial\,\beta}$$

Implying that the cost of learning is composed of two elements: the loss from not offering this contract to all consumers, given by $\sum_{\beta_j<\bar{\beta}} g(\beta_j)\,)\;\; E[\,W^*(\theta)\,|\,\beta_0\,]$; and the loss from facing consumers with beliefs different from the initial priors, given by $\sum_{\beta_j\geq\bar{\beta}} g(\beta_j)\,)\;\;(\beta_0 - \bar{\beta})\;\frac{\partial\,E[\,W^*(\theta)\,|\,\beta\,]}{\partial\,\beta}$. These terms coincide with the loses which drive the firm's trade-off (when choosing $\bar{\beta}$) between offering a contract to consumers with higher willingness to pay (i.e., higher beliefs) and offering a contract to a larger proportion of consumers.

Cost of preventing learning. This cost isolates the effect of price distortions by ignoring the effect of adverse selection. This is done by considering the case where the firm faces a single homogeneous belief β. Consider $\beta = \beta'$, such that we can write: $C_P(\vec{\beta}, \beta') \;\equiv\; \hat{\Pi}(\beta') - \hat{\Pi}(\beta'\,|\,\Theta)$, implying that this cost depends on the threshold belief served $\bar{\beta}$. Costs of distorting allocations have the same properties as before (by linearity of utility in tastes and quadratic costs of deviating from optimal quantities) implying that the cost of learning has very similar properties:

$$C_P(\,\vec{\Theta}\,|\,\bar{\beta}\,) \;\;=\;\; \sum_K \;\sigma^2(\Theta^K\,|\,\bar{\beta})\;\; Prob(\Theta^K)$$

Where this expression differs from the one developed before only by the fact that the conditional variances are calculated with the distribution functions of consumers

[12]This mean that, when consumers are uninformed, the new restriction to one contract is not binding and the firm generates profits given by $E[\,W^*(\theta)\,|\,\beta_0\,]$. On the other hand, when consumers are informed, the profits of a firm - which serves consumers above $\bar{\beta}$ - are given by $\sum_{\beta_j\geq\bar{\beta}} g(\beta_j)\,E[\,W^*(\theta)\,|\,\bar{\beta}\,]$, by maximising **??** above.

with belief β'.

Optimal contract and stationary profits. Under stationary contracts, as before, the total profits are:

Lemma 22. *Stationary profits, (in each period) when facing beliefs $\vec{\beta}$, choosing threshold $\bar{\beta}$ and offering pricing $\bar{\Theta}$, are given by:* $E[\,W^*(\theta)\,|\,\beta_0\,] \ - C_P(\,\bar{\Theta}\,|\,\bar{\beta}\,) \ - C_L(\,\vec{\beta}\,,\,\bar{\beta}\,)$

Importantly, it can be shown that, when the type heterogeneity is low, the expressions of the new cost of learning and the cost of preventing learning are going to be simplified, such that:

Lemma 23. *When type heterogeneity is small ($\alpha < \bar{\alpha}'$), the cost of learning and cost of preventing learning have the same properties as in unlimited number of contracts' case.*

This is not a trivial result, as the costs derived above depend on the threshold $\bar{\beta}$ which in turn is determined optimally by the overall belief vector $\vec{\beta}$ faced by the firm. However, this relationship simplifies considerably when type heterogeneity is small (i.e., α is small), as the firm serves all consumers.

When the firm decides which consumers to serve, it trades-off the benefit of excluding some consumers (i.e., being able to extract surpluses from remaining consumers with higher beliefs) and the implicit cost of the surplus foregone from excluded consumers. The former effect increases in how different beliefs are across consumers. However, when α decreases, the consumers find it harder to use past data to infer their types, meaning that difference between beliefs decreases. This means that, when the type heterogeneity α is small enough, the firm prefers to serve all consumers, such that the cost of learning simplifies to:

$$C_L(\,\vec{\beta}\,,\,\bar{\beta}\,) \ = \ (\beta_0 - \beta_1) \sum_i \left(W^*(\theta_i) - W^*(\theta_{i-1})\right) \left(F_L(\theta_i) - F_H(\theta_i)\right)$$

Implying that the cost of learning depends only on the lowest belief, from all the belief vector $\vec{\beta}$, and it is increasing in the taste level, which implies similar results to before:

Proposition 6. *When type heterogeneity is low enough ($\alpha < \bar{\alpha}'$), the optimal single contract, in a stationary contract is:*

- *3PT of size* 1*, whenever the taste level is low ($\theta_1 < \underline{\theta}_1'$)*

- *3PT of size* $1 < J < N$, *whenever the taste level is intermediate.* ($\underline{\theta}_1' < \theta_1 < \bar{\theta}_1'$)
- *3PT of size* N, *whenever the taste level is high* ($\bar{\theta}_1' < \theta_1$)

3.3.2 Two contracts

This case, where the firm can offer two contracts, combines features seen in the unlimited case (section 2) and the single contract case (section 3.1). The firm needs to decide which beliefs to target but, because it offers two contracts, it can partially screen consumers. When the firm faces a vector of beliefs $\vec{\beta}$ and chooses to offer contract to beliefs $\{\underline{\beta}$ and $\bar{\beta}\}$, the profits per period are given by:

$$\sum_{\underline{\beta}\leq\beta_j<\bar{\beta}} g(\beta_j)\ \Big(E[\,W(q_{\underline{\beta}}(\theta)\,,\,\theta)\,|\,\underline{\beta}\,]\ -\ IR(\underline{\beta})\Big)+\sum_{\beta_j\geq\bar{\beta}} g(\beta_j)\ \Big(E[\,W(q_{\bar{\beta}}(\theta)\,,\,\theta)\,|\,\bar{\beta}\,]\ -\ IR(\bar{\beta})\Big)$$

The difficulty is then to show how the firm is going to choose the thresholds to serve and how this new problem is going to affect the cost of learning and preventing learning. Fortunately, when type heterogeneity is small, the problem simplifies considerably, as we can show that:

Lemma 24. *When type heterogeneity is low ($\alpha < \bar{\alpha}''$), the firm decides to optimally serve all consumers, such that $\underline{\beta} = \beta_1$, and to offer $\bar{\beta} = \beta_0$*

This lemma relies on the concepts of cost of learning and cost of preventing learning, developed before, which are going to compose the profit function and can be used to find the optimal beliefs' thresholds. First, note that when the type heterogeneity is small and $\bar{\beta} \leq \beta_0$, profits per period are given by:

$$E[\,W^{*}(\theta)\,|\,\beta_0\,]\ -\ C_P(\bar{\Theta}|\underline{\beta})\sum_{\beta_j\leq\beta_0} g(\beta_j)\ -\ C_L(\bar{\beta},\underline{\beta},\vec{\beta})$$

Now, the cost of learning combines the three effects that decrease profits as discussed previously: (i) information rents from screening problem; (ii) profit loss from consumers excluded from the contract; (iii) loses in consumer's willingness to pay, as measured by how beliefs targeted differ from the prior belief without learning (β_0).

While the first effect increases with belief dispersion[13], the third effect increases with belief dispersion because the firm can tailor contracts better. Combining these two effects, however, leads to a negligible effect of belief dispersion on cost of learning. This means that the loses of profits from consumers excluded (the second effect), which depend mainly on the lowest belief, are going to dominate. As a result, the

[13]As implied by properties of costs developed in chapter 2.

firm is going to offer the lowest possible belief $\underline{\beta}$ possible and, hence, not exclude any consumers, meaning that the cost of learning is approximately given by:

$$(\beta_0 - \beta_1) \sum_{\beta_1 \leq \beta_j \leq 1} g(\beta_j) \, \frac{\partial E[\, W^*(\theta) \,|\, \beta \,]}{\partial \beta}$$

The absence of $\bar{\beta}$ from the cost of learning's expression, means that the upper threshold's decision is mainly determined by the cost of preventing learning. Then, the optimal $\bar{\beta}$ choice, which minimises C_P, is to set $\bar{\beta} = \beta_0$ by symmetry of probability distribution.

Hence, given that the cost of learning is the same as in the single contract and the cost of preventing learning follows a similar expression, these costs will have the same properties and the optimal contracts will have the same properties in terms of pricing:

Proposition 7. *When the firm is constrained to use two contracts and type heterogeneity is small ($\alpha < \bar{\alpha}''$), the optimal steady state contracts imply offering:*

i) Increasing pricing for both contracts, when θ_1 is low. ($\theta_1 < \underline{\theta}_1''$)

ii) Increasing pricing for lower contract and a 3PT for upper contract, when θ_1 is intermediate. ($\theta_1 \in (\underline{\theta}_1'', \bar{\theta}_1'')$)

iii)Increasing pricing for lower contract and a flat pricing for upper contract, when θ_1 is high. ($\bar{\theta}_1'' < \theta_1$)

3.4 3PT with competition

Augment the original model[14] by introducing a symmetric competitor to the (previously monopolistic) firm, and assume that the market of these (duopolistic) firms is horizontally differentiated. Label firms as L and R and, to model horizontal preferences, add brand-specific tastes $\epsilon(s)$ (for $s = L, R$) to the consumers' utility. This means that consumers are now characterised by how they value: the good sold (as captured by tastes θ); and the firms (or brands, captured by $\epsilon(L)$ and $\epsilon(R)$), such that utility is given by:[15]

$$\theta\, v(q) \;-\; P \;-\; \epsilon(s)$$

Each consumer's pair of brand-specific tastes $\epsilon(s)$ are: fixed over time; known to the consumer only; and drawn (at birth), independently for each consumer, from a joint uniform distribution over $[-K, K]^2$. Importantly, this means that the (privately known) brand-specific tastes $\epsilon(s)$ are independent from the consumers' types μ. We can represent firm L's relative preferences as a function of brand preferences:

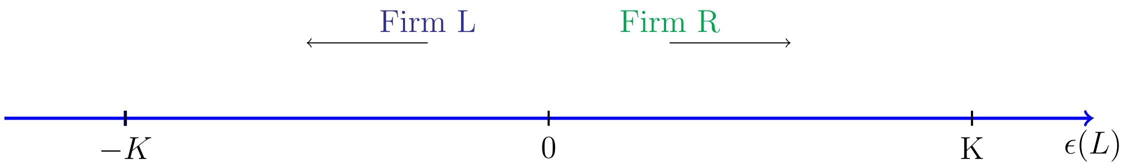

Figure 3.1: Brand preferences for firm L

These preferences differ from the Hotelling's horizontal differentiation and are closer to Rochet and Stole (2002)'s random outside options. The advantage of this kind of horizontal differentiated competition is that K can be truly interpreted as a measure of the market power. This can be contrasted with Hotelling's model, where the firm's profits are decreasing in the distance between firms[16], as consumers do not want to travel to firms that are too far away from their ideal points.

The firm's problem The firm is restricted to deterministic contracts, which implies that it can only screen consumers according to their beliefs and tastes (as before) and not according to their brand preferences[17]. Then, the only modification

[14] Keep assuming that each firm is identical to the monopolist described before: profits functions and contract sequential structure are identical.

[15] When the consumer buys, from firm $s \in \{L, R\}$, quantity q at cost P

[16] After a threshold where the competition softens.

[17] With deterministic mechanisms this is always true, as Rochet and Stole (2001) show.

to the firm's problem, as analysed before, is the participation constraint.[18] The proportion of consumers who buy from firm s, as implied by the firm's participation constraint, is determined by consumers who prefer brand s over both: the alternative $-s$; and an outside option with zero value. It is useful to write the expected utility in contract β by firm s as $u^s(\beta) \equiv E[\,U^s_\beta(\theta)\,|\,\beta\,]$. Then, we can see that for contracts designed for consumers with beliefs β the marginal consumer for firm s has brand taste $\epsilon_M(s)$:

$$\epsilon_M(s) \quad = \quad min\ \{\ u^s(\beta)\ \ ,\ \ u^s(\beta) - u^{-s}(\beta)\ \ + \epsilon(-s)\ \}$$

Before, for a large range of the firm's choices, the probability of participation would be largely constant, implying that the firm would always have an incentive to decrease surplus offered to consumers. Now, an important feature of this problem is that profits per period (from consumers with belief β_j) are given by:

$$\Pi(\beta_j) \quad = \quad g(\beta_j)\, prob\,(\,\epsilon \leq \epsilon_M(s)\,)\ \ [\,W(\beta_j)\ -\ u(\beta_j)\,]$$

Which means, as can be seen in the graph below, that the consumer's participation is a more smooth function of the surplus offered by the firm.

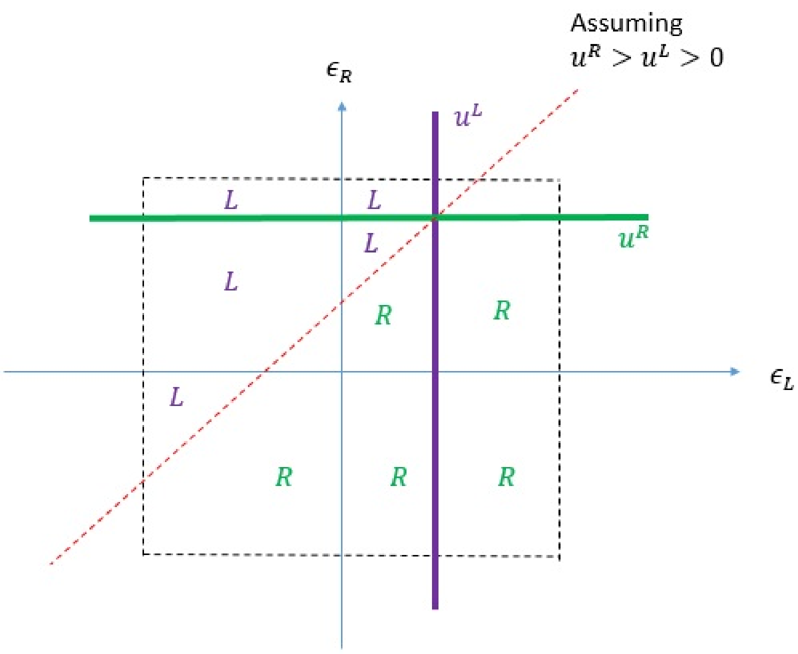

Figure 3.2: Consumers who buy from firm L and R for different values of $\epsilon(L)$ and $\epsilon(R)$ with asymmetric surpluses across firms.

[18]The participation constraint for firm S of a consumer with belief β and brand taste ϵ is:
$$PC_{\beta,S}\ \ :\ \ \ u^s(\beta)\ \ \geq\ \ min\,\{\,\epsilon(s)\ \ ,\ \ u^{-s}(\beta) + \epsilon(s) - \epsilon(-s)\,\}$$

Without adverse selection. I introduce this case to make the point that differentiated competition changes the firm's problem considerably. In particular, if adverse selection was absent - that is, the firm could offer exclusive contracts conditional on beliefs) - profits would increase with consumer learning.

Recall that, for the monopolist described until now, profits would be independent of belief dispersion/consumer learning if we could contract on consumer's beliefs[19], as profits in that case would be linear in beliefs.

Now, competition is going to change this relationship between profits and beliefs. With competition, the surplus offered to consumers decreases the direct revenues in each contract, as before, but it also increases the proportion of consumers who accept the contract. This implies that, even without adverse selection, it is not necessarily optimal to offer zero surpluses anymore. In particular, the optimal surplus is now given by:

$$u^s(\beta) \; = \; \frac{W^s(\beta)-K}{2} \; - \; \frac{(u^{-s}(\beta)+K)^2}{8K}$$

The fact that the firms' choice influences a new dimension - the proportion of consumers who are served - implies that their profits are going to be convex in beliefs. As shown below, this implies that the firm actually benefits from belief dispersion (and, hence, consumer learning):

Lemma 25. *Equilibrium competitive profits without adverse selection are increasing in the variance of consumers' beliefs.*

The important consequence of this lemma is that, if the firm could offer contracts conditional on beliefs in competitive markets, it would have no incentive to prevent consumer learning.

With adverse selection.[20] The costs of screening for the firm in this case are also going to increase a non-linear way, as compared to the monopolist case. In particular, in any symmetric equilibrium, per period profits are given by:

$$(\,1/2 \;\; - \;\; \frac{(\,K-u(\beta)\,)^2}{2}\,\frac{1}{4K^2}\,) \;\; [W(\beta)-u(\beta)]$$

[19]As seen in property 1 of C_L, the monopolist's profits were equal to the welfare generated by the contracts, as in this case the monopolist always chose to offer zero consumer surplus. This meant that the monopolist's expected profits were independent of consumer learning. The welfare of these contracts is linear in beliefs, which in turn are equal to the prior belief β_0, independently of consumer learning, by the Martingale property.

[20]Equilibrium existence is guaranteed by similar arguments to the one developed in Rochet and Stole (2002), page 290, and Economides (1989).

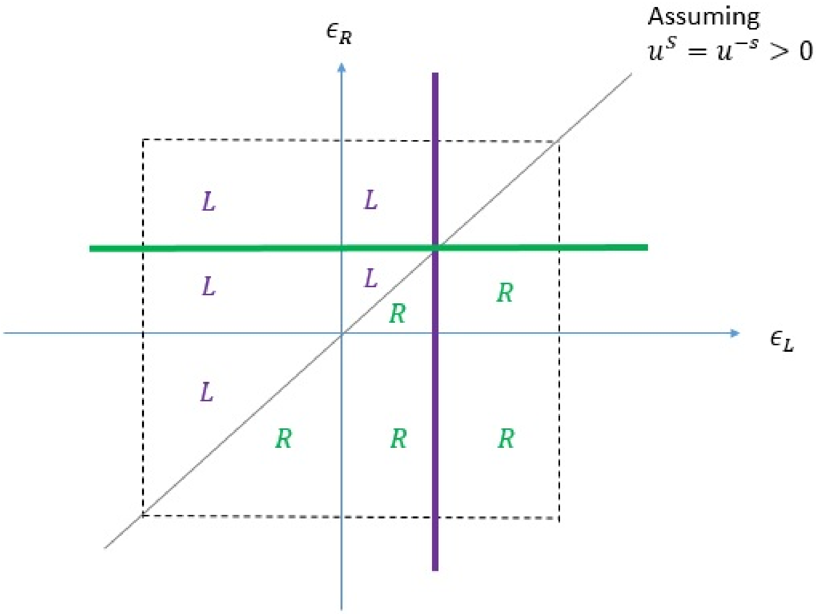

Figure 3.3: Consumers who buy from firm L and R for different values of $\epsilon(L)$ and $\epsilon(R)$ with symmetric surpluses across firms.

When $u^s(\beta) = u^{-s}(\beta) = u(\beta)$. Moreover, even though the surplus' levels are different, the screening problem does not change. ICs are still binding in its downward local version:[21]

Lemma 26. *Local (downward) ex ante incentive constraints are going to be binding, as before.*

Finally, we can show that, when the firm has sufficient market power, it will choose to offer contracts with the same pricing structure as described in proposition 4. Write the firms profits from each belief β_j:

$$[\ 1/2 \ + \ u^s(\beta_j)\ 2K \frac{1}{4K^2} \ - \ \frac{(\,K + u^{-s}(\beta)\,)^2}{2} \frac{1}{4K^2}\ ]\ [W^s - u^s(\beta_j)]$$

We can see that the marginal impact of surpluses on the probability of participation is decreasing in the level of K, while its effect on profits per unit sold are identical to before.[22] This argument leads to the following proposition:

[21] While brand preferences offer a new motive to offer positive consumers surplus implying that the ICs could potentially be relaxed, this is not the case in our model for similar reasons as developed by Rochet and Stole (2002).

[22] Where, formally, we need to use that the surpluses are bounded to make this argument.

Proposition 8. *For small type heterogeneity and large enough market power ($K > \bar{K}$), duopolists will offer 3PT under similar conditions as monopolists.*

Where the proof takes advantage of the simple fact that the profits of the firm are going to be less influenced by the consumers who are marginally excluded by increasing the surplus extraction the higher is the firm's market power. When the firm takes the probability of participation as largely constant, it will optimise profits as a monopolist would, leading to the pricing result developed in chapter 2.

3.5 Discussion

The model developed in chapter 2 describes a new mechanism which explains why firms distort pricing - by offering 3PT - in subscription contracts. In chapter 3, I have shown that, under some conditions, this mechanism is robust to considerations that are relevant in these markets such as limited contracts and competition. Here I discuss the role of the simplifying assumptions needed to get these results and further issues.

Firstly, assuming that consumers are myopic could be relaxed while keeping the results qualitatively similar.[23] Far-sighted consumers would not be able to (individually) influence future contracts with their decisions as there is a continuum of consumers, so these issues can be ruled out in equilibrium.[24] However, in this case, consumers with taste θ could try to use their consumption decision $q_\beta(\theta')$ (or taste report) to communicate information about their past taste θ to their future self. This means that the firm would need to satisfy a more complex ex post incentive compatibility constraint.

However, this constraint would not bind for all consumers. Those consumers who had tastes realisations in regions where marginal pricing is positive will not engage in this deviation, as they already going to be able to remember their tastes in the future. On the other hand, consumers with taste realisations in regions of flat pricing[25] could engage in this deviation by misreporting their type today to communicate it to their future self.[26]

Then, for consumers with tastes in a region of flat prices, only the consumers with the lowest tastes really have an incentive to misreport their taste. Only these consumers, by recalling their past taste, can avoid accepting in the next period a contract which has a negative expected value[27].

Finally, I can show that in 3PT, the consumers in the only flat region can only misreport their taste by pretending to have a higher taste. This means that the consumers who have the highest incetive to deviate are also the ones who suffer the most from it. Pretending to have a taste considerably higher will decrease

[23]Relaxing this assumption and introducing cost of processing information would make the consumers in this model rationally inattentive (Gabaix, 2014).

[24]Dynamic smoothing consumption considerations are ignored by assuming that the money spent in this market is a small fraction of the consumer's income.

[25]Such that their taste θ is offered the same price as another taste $\theta' \neq \theta$: $P_\beta(\theta) = P_\beta(\theta')$

[26]In a multi-self equilibrium as considered by Piccione and Rubinstein (1997).

[27]Remember that consumers of a flat region will remember only that their taste is in that range and will be offered a contract in the next period that charges them their expected surplus (subtracted of information rents).

their payoff today considerably whenever static incentive compatibilities hold. This implies that, under some conditions, this new constraint is more likely to be binding for other pricing structures than to 3PT, meaning that the results developed here will not change.[28].

Secondly, the 3PT as modelled here differ from the ones observed in markets in two ways: (1) The pricing is flat/constant in tastes but not quantities[29] ; (2) The pricing is not technically above marginal cost for large tastes.

The first issue is easy to address as shown in Grubb (2001). This problem can be solved by assuming that consumers have satiated preferences, which creates a link between zero marginal pricing in tastes and in quantities. Under this modification, the firm will offer zero marginal pricing over quantities, under similar conditions to the ones developed in the main text. To address the second issue, the model would have to be augmented to include other realistic features of learning as identified by Ater and Landsman (2016). When learning is not fully rationa such that consumers over-infer from their small sample and the intensity of this memory is related to the amount of their previous bill size, the firm will also have incentives to increase prices for large usage to manipulate consumers' beliefs.

Lastly, there a few technical assumptions that play an important role. Low type heterogeneity is a condition that is technically important for the problem to be tractable. Without this assumption, other issues about the distribution of beliefs will play a role, leading to a more complicated relationship between consumer learning and profits. Assuming that consumers live for two periods, implies that consumers never fully learn their type. As long as consumers information depreciates over time or there can be systemic changes in the properties in these markets (e.g., technological innovation) this assumption seems plausible. On the other hand, this also suggests that maybe this pricing will be used less often in more mature markets, which seems to be consistent with evidence from the cellphone service industry[30].

[28]Other than the reducing the size of implementable 3PT.

[29]In the present model, flat pricing in a contract (that is, $P_\beta(\theta) = P_\beta(\theta')$) is going to imply that the consumers with these tastes are supplied the same quantity. This is not equivalent to zero marginal pricing in quantities (that is, $\frac{\partial P_\beta(.)}{\partial q} = 0$).

[30]But it is hard to argue that this link is casual as other technological changes occur normally as the industries mature.

3.6 Appendix

3.6.1 Appendix of section 3.2

Proofs in this section are direct implications of the the fact that, when the two facts described below hold, the problem in the OLG model is equivalent to the two period problem (when δ is close to one) meaning that the optimal pricing will be the same.

Fact 1. Problem can be solved in two steps. Ignore pricing constraints (which are implied by decision to prevent learning) and write the expression for virtual surplus[31]:

$$\Phi_j \equiv \; E[\,W(q\,,\,\theta_i - X_i(\beta_j))\,|\,\beta_j\,] \qquad \text{and normalised information rents are} \qquad X_i(\beta_j)$$

Then, incorporate the pricing constraints. This will reduce profits in the similar way to before, giving rise to the same cost of preventing learning which now affect normalised tastes (that is, $\hat{\theta}_i(\beta) = \theta_i - X_i(\beta)$) implied in the virtual surplus solution. This implies that profits can be written as:

$$\Pi(\vec{\beta}) \;\; = \;\; \hat{\Pi}(\beta_0) \;\; - \;\; C_L(\vec{\beta_t}) - \;\; \hat{C}_P(\bar{\Theta}_t) \;\; g(\beta_Y)$$

Where, whenever the monotonicity constraints are not violated, $C_L(\vec{\beta})$ is as defined before, and $\hat{C}_P \;=\; \sigma^2(\hat{\Theta}_k\,|\,\beta_Y)$ and $\hat{\Theta}_k$ are transformed sub-set of tastes prevented from learning. That is, tastes are transformed to include normalised info rents.

Fact 2. As before, the limit of normalised information rents as type heterogeneity goes to zero is zero. This means that profits in this section are going to approach profits of the previous chapter and for type heterogeneity sufficiently small, the same results will hold.

An implication of this lemma is that the firm's optimal pricing in each period is independent of the current beliefs. In period t, the firm decides on which tastes to prevent learning (in contract offered to young consumers), formally written as pricing $\bar{\Theta}_t$, to minimise the following expression:

$$-\hat{C}_P(\bar{\Theta}_t) - \delta \; C_L(\vec{\beta_{t+1}})$$

[31]Since as before since information rents will be determined in the same way and bind.

Where $\vec{\beta}_{t+1}$ is the future belief vector implied by pricing $\bar{\Theta}_t$. The solution to this problem in period t does not depend on current beliefs $\vec{\beta}_t$ as long as none of the factors at stake depend on current beliefs, which is true when $\hat{C}_P(\bar{\Theta})$ is close to $C_P(\bar{\Theta})$. This last statement is true when α is small, such that the normalised information rents are small. Whenever δ is close to one, this means that focusing on steady state contracts is wlog.

3.6.2 Appendix of 3.3.1

Proof of Lemma 5 and 6.

Proof. 1) *Cost of learning.* When the firm serves consumers with beliefs equal or above $\bar{\beta}$:

$$C_L(\vec{\beta}, \bar{\beta}) \;=\; \textstyle\sum_{\beta_j<\bar{\beta}} g(\beta_j)\,)\, E[\,W^*(\theta)\,|\,\beta_0\,] + \sum_{\beta_j\geq\bar{\beta}} g(\beta_j)\,)\,(\beta_0-\bar{\beta})\,\frac{\partial E[\,W^*(\theta)\,|\,\beta\,]}{\partial\beta}$$

Derive this expression by understanding that profits decrease with learning because:

i) The firm can only charge the willingness to pay of the marginal consumer. Profits from a single contract are proportional to: $E[W^*(\theta)\,|\,\beta\,] \;=\; E[W^*(\theta)\,|\,\beta_0\,] - (\,\beta_0-\beta\,)\,\frac{\partial E[\,W^*(\theta)\,|\,\beta\,]}{\partial\beta}$

ii) The firm does not sell to all consumers, which decreases the profits by the proportion of consumers not buying the good multiplied by the potential profits: $\sum_{\beta_j<\bar{\beta}} g(\beta_j)\,)\, E[\,W^*(\theta)\,|\,\beta_0\,]$.

What is the optimal $\bar{\beta}$? As α goes down, the difference between any two beliefs $|\beta_{j+1}-\beta_j|$ decreases. Implying that the benefit of facing consumers with higher beliefs decreases, leading the firm to sell to all consumers: $\bar{\beta}=\beta_1$.

2) *Cost of preventing learning.* Using that, when the firm does not prevents learning, it can extract all the welfare from this consumer and get profits equal to $E[\,W^*(\theta)\,|\,\beta\,]$; and when the firm prevents learning $\hat{\Pi}(\beta|\Theta_k) \;=\; E[\,W^*(\hat{\theta})\,|\,\beta\,]$ Where $\hat{\theta}=\theta$ when $\theta\notin\Theta^K$ and $\hat{\theta}=\mu^K$ when $\theta\in\Theta^K$. For small enough α, we can consider that the cost of preventing learning satisfies $C_P(\bar{\Theta}|\bar{\beta}) \approx C_P(\bar{\Theta}|\beta_0) \equiv C_P(\bar{\Theta})$, as beliefs converge as α decrease.

3) *Summarising:* the cost of preventing learning is the same as before and the new cost of learning shares the same properties as before. When the firm serves all consumers, the cost of learning can be simplified to:

$$C_L(\,\vec{\beta}\,|\,\bar{\beta}\,) \;=\; (\beta_0-\beta_1)\sum_i (W^*(\theta_i)-W^*(\theta_{i-1}))\,(F_L(\theta_i)-F_H(\theta_i))$$

Where it is easy to see that the cost of learning, from the whole belief vector $\vec{\beta}$:

- Only depends on the lowest belief β_1 and the prior belief β_0.
- Increases with taste level θ_1 and with type heterogeneity α. □

Proof of Proposition 2.

Same argument as before, since cost of learning and cost of preventing learning have the same properties.

3.6.3 Appendix of 3.3.2

Proof of Lemma 7.

Step 1. *Incorporating constraints into profit function.* The profits when the firm chooses to tailor the contract to beliefs $\{\underline{\beta}\,,\,\bar{\beta}\}$ are given by:

$$\sum_{\underline{\beta}\le\beta_j<\bar{\beta}} g(\beta_j)\ \Big(\ E[\,W(q_{\underline{\beta}},\theta)\,|\,\underline{\beta}\,]\ -\ IR(\underline{\beta})\ \Big) + \sum_{\beta_j\ge\bar{\beta}} g(\beta_j)\ \Big(\ E[\,W(q_{\bar{\beta}},\theta)\,|\,\bar{\beta}\,]\ -\ IR(\bar{\beta})\ \Big)$$

a) Following the same steps as in the unlimited contracts (section 2), we can simplify the information rents and show that: information rents for lowest belief are zero: $IR(\underline{\beta}) = 0$; and information rents for highest belief are pinned down by downward ex ante IC:

$$IR(\bar{\beta})\ =\ \textstyle\sum_i\,(F_L(\theta_i) - F_H(\theta_i))\,\Delta\,v(q_{\underline{\beta}}(\theta_{i-1})\,(\,\bar{\beta} - \underline{\beta}\,).$$

b) Which consumers are affected by Price distortions? There are two cases, depending on which contract serves the young consumers, given by $\bar{\beta} > \beta_0$ and $\bar{\beta} \le \beta_0$. For concreteness, consider for now the case where $\bar{\beta} \le \beta_Y$ to rewrite profits as:

$$\sum_{\underline{\beta}\le\beta_j<\bar{\beta}} g(\beta_j)\ \ E[\,W(q_{\underline{\beta}}\,,\,\theta - X(\underline{\beta})\,)\,|\,\underline{\beta}\,]\ +\ \sum_{\beta_j\ge\bar{\beta}} g(\beta_j)\ \ E[\,W(\bar{q}_{\bar{\beta}}\,,\,\theta)\,|\,\bar{\beta}\,]$$

When the firm pools quantities for contract $\bar{\beta}$. Substituting optimal quantities and using information rents above, we can write profits as a function of welfare with virtual surplus:

$$\sum_{\underline{\beta}\le\beta_j<\bar{\beta}} g(\beta_j)\ \ E[\,W^*(\,\theta - X(\underline{\beta})\,)\,|\,\underline{\beta}\,]\ +\ \sum_{\beta_j\ge\bar{\beta}} g(\beta_j)\ \ \big(\,E[\,W^*(\,\theta\,)\,|\,\bar{\beta}\,]\ -\ C_P(\bar{\Theta}|\bar{\beta})\,\big)$$

Using the fact that pooling quantities is going to have the same effect on profits as before.

STEP 2. *Finding the cost of learning.* Here we show that the profits are going to be written as:

$$E[\, W^*(\,\theta\,)\,|\,\beta_0\,] \;\; - \;\; C_P \sum_{\underline{\beta} \leq \beta_j < \bar{\beta}} g(\beta_j) \;\; - \;\; C_L^{Overall}$$

Here, we show that the new cost of learning is going to have three elements: distortions from information rents; distortions from beliefs below priors; and distortions from excluded consumers. We analyse each in turn:

a) **Screening.** The cost of learning from screening follows the familiar formula from the case in section 2 and can again be separated from the profits:

$$\sum_{\underline{\beta} \leq \beta_j < \bar{\beta}} g(\beta_j) \;\; \Big(\, E[\, W^*(\,\theta\,)\,|\,\underline{\beta}\,] \; - C_L(\vec{\beta}) \Big) \; + \sum_{\beta_j \geq \bar{\beta}} g(\beta_j) \;\; \big(\, E[\, W^*(\,\theta\,)\,|\,\bar{\beta}\,] \; - \; C_P \, \big)$$

Where the normalised information rents are proportional to the difference between the beliefs in the two contracts:

$$X_i(\underline{\beta}) \;\; = \;\; (\, \bar{\beta} - \underline{\beta} \,) \;\; (F_L(\theta_{i+1}) - F_H(\theta_{i+1})) \;\; \frac{\sum_{\beta_j \geq \bar{\beta}} g(\beta_j)}{\sum_{\underline{\beta} \leq \beta_j < \bar{\beta}} g(\beta_j) \;\; f_{\underline{\beta}(\theta_i)}}$$

And the cost of learning $C_L(\vec{\beta})$ has the familiar expression: $\sum_i \; X_i(\underline{\beta}) \; (\, 2\theta_i - X_i(\underline{\beta}) \,) \;\; f_{\underline{\beta}}(\theta_i) \; \sum_{\underline{\beta} \leq \beta_j < \bar{\beta}} g(\beta_j) \,)$, which makes clear the dependency of these costs on the normalised information rents.

b) **Belief choice in limited contracts.** The cost of facing lower beliefs than the uninformed ones:

$$E[W^*(\theta)\,|\,\beta\,] \;\; = \;\; E[W^*(\theta)\,|\,\beta_0\,] \; - \; (\, \beta_0 - \beta \,) \; \textstyle\sum_i W^*(\theta_i) \; (f_H(\theta_i) - f_L(\theta_i))$$

It is going to be useful to define $L \;=\; \sum_i W^*(\theta_i) \; (f_H(\theta_i) - f_L(\theta_i))$ and rewrite this term using increments in welfare (as we do for information rents' formula):

$$L \;=\; \sum_i \; (\, W^*(\theta_i) - W^*(\theta_{i-1}) \,) \, (\, F_L(\theta_i) - F_H(\theta_i) \,) \;\;\; = \;\; \sum_i \; (\, 2\Delta(\, \theta_{i-1} + \frac{\Delta}{2} \,) \,) \, (\, F_L(\theta_i) - F_H(\theta_i) \,)$$

Using (in the last equality) that $W^*(\theta_i) - W^*(\theta_{i-1})$ is equal to: $2\Delta \, (\, \theta_{i-1} + \frac{\Delta}{2} \, , \,)$, given the functional form assumed.

c) **The cost of excluding consumers.** This contract implies that consumers below $\underline{\beta}$ are excluded from the contract. The loses in profits when this happens are given by:

$$\textstyle\sum_{\beta_j<\underline{\beta}} g(\beta_j)\ E[\,W^*(\theta)\,|\,\beta_0\,].$$

d) **Combining all terms of the cost of learning.** Using previous steps, $C_L^{Overall}$ can be written as (in two steps):

$$(\,\beta_0 - \underline{\beta}\,) \sum_{\underline{\beta}\le\beta_j\le 1} g(\beta_j)\ L\ +\ \underbrace{(\,\bar{\beta} - \underline{\beta}\,) \sum_{\beta_j\ge\bar{\beta}} g(\beta_j)\ [G - L\,]}_{C_L(\vec{\beta})\ -\ (\bar{\beta}-\underline{\beta})\ \sum_{\beta_j\ge\bar{\beta}}\ g(\beta_j)\ L}\ + \sum_{\beta_j<\underline{\beta}} g(\beta_j)\ E[\,W^*(\theta)\,|\,\beta_0\,]$$

Where G is the term of X that is multiplied by $(\,\bar{\beta} - \underline{\beta}\,)\ \sum_{\beta_j\ge\bar{\beta}} g(\beta_j)$, such that the second term above can be simplified to:

$$\underbrace{(\,\bar{\beta} - \underline{\beta}\,) \sum_{\beta_j\ge\bar{\beta}} g(\beta_j) \left[\sum_i \Delta\,(F_L^{i+1} - F_H^{i+1})\,(\,2\theta_i - X_i(\underline{\beta})\ - 2(\,\theta_i + \Delta/2\,)\right]}_{C_L(\vec{\beta})\ -\ (\bar{\beta}-\underline{\beta})\ \sum_{\beta_j\ge\bar{\beta}}\ g(\beta_j)\ L}$$

And cancelling common terms to get:

$$-\ (\,\bar{\beta} - \underline{\beta}\,) \sum_{\beta_j\ge\bar{\beta}} g(\beta_j) \left[\sum_i \Delta\,(F_L^{i+1} - F_H^{i+1})\,(\,X_i(\underline{\beta}) + \Delta/2\,)\right]$$

Implying that the overall cost of learning becomes:

$$(\,\beta_0 - \underline{\beta}\,)\ \textstyle\sum_{\underline{\beta}\le\beta_j\le 1}\ g(\beta_j)\ L\ \underbrace{-\ (\,\bar{\beta} - \underline{\beta}\,) \displaystyle\sum_{\beta_j\ge\bar{\beta}} g(\beta_j) \left[\sum_i \Delta\,(F_L^{i+1} - F_H^{i+1})\,(\,X_i(\underline{\beta}) + \Delta/2\,)\right]}_{C_L(\vec{\beta})\ -\ (\bar{\beta}-\underline{\beta})\ \sum_{\beta_j\ge\bar{\beta}}}\ + \textstyle\sum_{\beta_j<\underline{\beta}}\ g(\beta_j)\ E[\,W^*(\theta)\,|\,\beta_0\,]$$

STEP 3. *Finding optimal thresholds.* Given some belief vector $\vec{\beta}$, what is the optimal choice of $\underline{\beta}$ and $\bar{\beta}$?

a) **Solve for the lower bound $\underline{\beta}$.** For the same argument as before (if α is small, the effect of threshold is small on beliefs and normal in terms of proportion of consumers excluded), it is optimal to set $\underline{\beta} = \beta_1$ to minimize cost of excluding consumers. Mathematically, this argument relies on the fact that $\underline{\beta}$ influences only these two terms:

$$(\beta_0 - \underline{\beta}) \sum_{\underline{\beta} \leq \beta_j \leq 1} g(\beta_j) \, L \quad + \quad \sum_{\beta_j < \underline{\beta}} g(\beta_j) \; E[\, W^*(\theta) \,|\, \beta_0 \,]$$

And L can be made small by decreasing α.

b) **Consider upper bound $\bar{\boldsymbol{\beta}}$.** This threshold $\bar{\beta}$ influences both the cost of learning and the cost of preventing learning. The terms of these cost terms relevant for this decision are:

$$- \; (\bar{\beta} - \underline{\beta}) \sum_{\beta_j \geq \bar{\beta}} g(\beta_j) \left[\sum_i \Delta \, (F_L^{i+1} - F_H^{i+1}) \, (\, X_i(\underline{\beta}) + \Delta/2 \,) \right] \; + \; \sum_{\cdots} g(\beta_j) \; C_P$$

Where the second term depends on the choice of $\bar{\beta}$ such that it can be either: $\sum_{\beta_j \geq \bar{\beta}} \, g(\beta_j) \; C_P$ for $\bar{\beta} \leq \beta_Y$ or $\sum_{\underline{\beta} \leq \beta_j < \bar{\beta}} \, g(\beta_j) \; C_P$ for $\bar{\beta} > \beta_Y$.

Think about how the solution affects each term. First, to minimise the first term of this expression, the optimal solution would be close to $\frac{\underline{\beta} + \beta^{max}}{2}$ which is $> 1/2$ if firm has flat region at the beginning. Second, the C_P term is minimised by setting $\bar{\beta} = \beta_0$ since symmetry that probability is 1/2 at that point and larger at any other point. Then, when α is small the term that multiplies the first expression is small and solution will be $\bar{\beta} = \beta_0$, taking advantage of discreteness.

Proof of Proposition 3.

Same argument as before, since cost of learning and cost of preventing learning have the same properties.

3.6.4 Appendix of section 3.4

Proof of lemma 8.

The marginal consumer (with belief β) - indifferent between accepting firm s's contract or not has taste benefits given by:

$$\epsilon(s) \;=\; min\;\{\; u^s(\beta)\;\;,\;\; u^s(\beta) - u^{-s}(\beta) \;\; + \epsilon(-s)\;\}$$

Meaning that the profits from belief β_j are given by:

$$g(\beta_j)\;\; Prob(\,\epsilon < \epsilon_M(s)\,)\;[\,W^s(\beta_j) - u^s(\beta_j)\,]$$

Given that we can analyse the case of each belief separately, drop the β's dependency for simplicity. Then, using figure 2, it can be shown that $Prob(\,\epsilon < \epsilon_M(s)\,)$ is equal to (for the case $u^s \geq u^{-s}$)[32]

$$\underbrace{1/2}_{\text{Participation proportion with zero surplus}} \quad + \quad \underbrace{u^s\;\; 2K\,\frac{1}{4K^2}}_{\text{Participation gains/loses from higher/lower surplus}} \quad - \quad \underbrace{\frac{(\,K + u^{-s}\,)^2}{2}\;\frac{1}{4K^2}}_{\text{Consumers lost to the other firm}}$$

Solve interior solution using the FOC, which gives you a best response for firm s (for $u^s \geq u^{-s}$):

$$u^s \;=\; \frac{W^s - K}{2} \;-\; \frac{(\,u^{-s} + K\,)^2}{8K}$$

Where interior symmetric solution, using symmetry (and dropping firm specific upperscripts), becomes:

$$u^* \;=\; -5K \;+\; \sqrt{\,25K^2 \;+\; 5K^2 \;-\; W}$$

Which gives us equilibrium profits equal to:

$$(6K^2 - \sqrt{\,25K^2 \;+\; 5K^2 \;-\; W})\;[\;(6K^2 - \sqrt{\,25K^2 \;+\; 5K^2 \;-\; W})\;+\;(W - K)\;]$$

Which are increasing and convex in welfare and, hence, convex in beliefs.

Proof of lemma 9.

[32]For the opposite case proportion needs to be added a correction term $-\frac{(u^s - u^{-s})^2}{2}\;\frac{1}{4\,K^2}$

The full problem can be written as maximising:

Problem 3.

$$\sum_j g(\beta_j)\, Prob(\,\epsilon < \epsilon_M(s)\,)\,[\,W(\beta_j) - u(\beta_j)\,]$$

Subject to

$$\textit{Ex post IC} \qquad \textit{(As before)}$$

$$\textit{Ex ante IC:} \qquad u_{j-1} + T_{j-1} \;\leq\; u_j \;\leq\; u_{j-1} + T_j$$

Where $\quad T_j(q_j) \;\equiv\; \sum_{k=1}^{N} \Delta\, q_j^i\, (F_L^i - F_H^i)\,(\beta_j - \beta_{j-1})$

Where it is clear that:

- If ex ante IC is not binding, then solution would be given by u^* and, hence, efficient quantities q_i^* (as profits are increasing in welfare in this case).
- However, substituting in the ex ante this solution shows that the solution violates this constraint, as

Proof of Proposition 4.

a) The probability of participation tends to a constant as K grows large.

The profits in any symmetric equilibrium are:

$$(\,1/2 \quad - \quad \frac{(\,K - u\,)^2}{2}\,\frac{1}{4K^2}\,)\,[W - u]$$

- The surplus that the firm can offer u is bounded between $-K$ and $min\{K, W\}$.

- For a fixed surplus u, the probability of participation tend to profits with exogenous probability of participation as K grows.

- For interior solution u^*, previously found, the optimal probability of participation converges to a constant as K increases.

- For corner solutions $u = -K$ the same is true.

b) The first order conditions tend to the ones of monopolist, as K increases.

Incentives for u are going to converge to the case with exogenous probability of participation, as K grows, since profits are:

$$[\;\; 1/2 \;\;+\;\; u^s(\beta)\;\; 2K\,\frac{1}{4K^2} \;\;-\;\; \frac{(\,K + u^{-s}(\beta)\,)^2}{2}\,\frac{1}{4K^2}\,]\,[W^s - u^s(\beta)]$$

Bibliography

Agarwal, Sumit, et al. Learning in the credit card market. No. w13822. National Bureau of Economic Research, 2008.

Ali, S. Nageeb. "Learning self-control." The Quarterly Journal of Economics 126.2 (2011): 857-893.

Arrow, Kenneth J., and Partha S. Dasgupta. "Conspicuous consumption, inconspicuous leisure." The Economic Journal 119.541 (2009): F497-F516.

Ater, Itai, and Vardit Landsman. "Do customers learn from experience? Evidence from retail banking." Management Science 59.9 (2013): 2019-2035.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. "Trends in US wage inequality: Revising the revisionists." The Review of economics and statistics 90.2 (2008): 300-323.

Bagwell, Laurie Simon, and B. Douglas Bernheim. "Veblen effects in a theory of conspicuous consumption." The American Economic Review (1996): 349-373.

Becker, Gary S., Kevin M. Murphy, and Edward Glaeser. "Social Markets and the Escalation of Quality: The World of Veblen Revisited." Social Economics: Market Behavior in a Social Environment (2000): 84-104.

Bénabou, Roland. "The Economics of Motivated Beliefs." Revue d'économie politique 125.5 (2015): 665-685.

Besley, Timothy, and Anne Case. Diffusion as a learning process: Evidence from HYV cotton. No. 228. 1994.

Bloch, Francis, Vijayendra Rao, and Sonalde Desai. "Wedding celebrations as conspicuous consumption signaling social status in rural India." Journal of Human Resources 39.3 (2004): 675-695.

Bolton, Patrick, and Christopher Harris. "Strategic experimentation." Econometrica 67.2 (1999): 349-374.

Bergemann, Dirk, and Juuso Valimaki. "Bandit problems." (2006).

Bloom, Nicholas, et al. "Management practices across firms and countries." The Academy of Management Perspectives 26.1 (2012): 12-33.

Carruthers, Peter. The opacity of mind: an integrative theory of self-knowledge. OUP Oxford, 2011.

Charles, K. K., E. Hurst, and N. Roussanov. "Conspicuous consumption and race." Quarterly Journal of Economics." (2008).

Clark, Andrew E., and Andrew J. Oswald. "Satisfaction and comparison income." Journal of public economics 61.3 (1996): 359-381.

Courty, Pascal, and Li Hao. "Sequential screening." The Review of Economic Studies 67.4 (2000): 697-717.

Dow, James. "Search decisions with limited memory." The Review of Economic Studies 58.1 (1991): 1-14.

Duesenberry, James Stemble. "Income, saving, and the theory of consumer behavior." (1949).

Dupas, Pascaline. "Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment." Econometrica 82.1 (2014): 197-228.

Economides, N., 1989, "Symmetric Equilibrium Existence and Optimality in Differentiated Product Markets", Journal of Economic Theory, Vol. 47, 178-194.

Feder, Gershon, Richard E. Just, and David Zilberman. "Adoption of agricultural innovations in developing countries: A survey." Economic development and cultural change 33.2 (1985): 255-298.

Ferrer-i-Carbonell, Ada. "Income and well-being: an empirical analysis of the comparison income effect." Journal of Public Economics 89.5 (2005): 997-1019.

Foster, Andrew D., and Mark R. Rosenzweig. "Learning by doing and learning from others: Human capital and technical change in agriculture." Journal of political Economy (1995): 1176-1209.

Gabaix, Xavier. A sparsity-based model of bounded rationality. No. w16911. National Bureau of Economic Research, 2011.

Ghatak, Maitreesh. "Theories of poverty traps and anti-poverty policies." The World Bank Economic Review 29.suppl 1 (2015): S77-S105.

Glaeser, Edward L., et al. "Measuring trust." Quarterly Journal of Economics (2000): 811-846.

Grubb, Michael D. "Selling to overconfident consumers." The American Economic Review 99.5 (2009): 1770-1807.

Grubb, Michael D., and Osborne, Matthew. "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock." American Economic Review 105.1 (2015): 234-271.

Goettler, Ronald L., and Karen Clay. "Tariff choice with consumer learning and switching costs." Journal of Marketing research 48.4 (2011): 633-652.

Harris, Milton, and Robert M. Townsend. "Resource allocation under asymmetric information." Econometrica: Journal of the Econometric Society (1981): 33-64.

Heffetz, Ori, and Robert H. Frank. "Preferences for status: Evidence and economic implications." HANDBOOK OF SOCIAL ECONOMICS, Jess Benhabib, Alberto Bisin, Matthew Jackson, eds 1 (2008): 69-91.

Hirshleifer, David, and Ivo Welch. An economic approach to the psychology of change: Amnesia, inertia. and impulsiveness. Working Paper, 2001.

Johnson, Justin P., and David P. Myatt. "On the simple economics of advertising, marketing, and product design." The American Economic Review 96.3 (2006): 756-784.

Kahneman, Daniel. Thinking, fast and slow. Macmillan, 2011.

Keller, Godfrey, and Sven Rady. "Strategic experimentation with Poisson bandits." Theoretical Economics 5.2 (2010): 275-311.

Laffont, Jean-Jacques, and Jean Tirole. "Using cost observation to regulate firms." The Journal of Political Economy (1986): 614-641.

Lambrecht, Anja, and Bernd Skiera. "Paying too much and being happy about it: Existence, causes, and consequences of tariff-choice biases." Journal of marketing Research 43.2 (2006): 212-223.

Lazear, Edward P., and Sherwin Rosen. "Rank-Order Tournaments as Optimum Labor Contracts." Journal of Political Economy 89.5 (1981): 841-864.

McLaughlin, Kenneth J. "Aspects of tournament models: A survey." Research in labor economics 9.1 (1988): 225-56.

Miravete, Eugenio J., and Ignacio Palacios-Huerta. "Consumer inertia, choice dependence, and learning from experience in a repeated decision problem." Review of Economics and Statistics 96.3 (2014): 524-537.

Mirrlees, James A. "An exploration in the theory of optimum income taxation." The review of economic studies 38.2 (1971): 175-208.

Moav, Omer, and Zvika Neeman. "Status and poverty." Journal of the European Economic Association 8.2-3 (2010): 413-420.

Mullainathan, Sendhil. "A memory-based model of bounded rationality." Quarterly Journal of Economics (2002): 735-774.

Myerson, Roger B. "Incentive compatibility and the bargaining problem." Econometrica: journal of the Econometric Society (1979): 61-73.

Nelson, Richard R., and Edmund S. Phelps. "Investment in humans, technological diffusion, and economic growth." The American economic review 56.1/2 (1966): 69-75.

O'Donoghue, Ted, and Matthew Rabin. "The economics of immediate gratification." Journal of Behavioral Decision Making 13.2 (2000): 233.

Parente, Stephen L., and Edward C. Prescott. "Barriers to technology adoption and development." Journal of political Economy (1994): 298-321.

Piccione, Michele, and Ariel Rubinstein. "On the interpretation of decision problems with imperfect recall." Games and Economic Behavior 20.1 (1997): 3-24.

Piketty, Thomas, and Emmanuel Saez. The evolution of top incomes: a historical and international perspective. No. w11955. National Bureau of Economic Research, 2006.

Ray, Debraj, and Arthur Robson. "Status, Intertemporal Choice, and Risk-Taking." Econometrica 80.4 (2012): 1505-1531.

Rochet, Jean-Charles, and Lars A. Stole. "Nonlinear pricing with random participation." The Review of Economic Studies 69.1 (2002): 277-311.

Rochet, Jean-Charles, and Lars A. Stole. "The economics of multidimensional screening." Econometric Society Monographs 35 (2003): 150-197.

Rothschild, Michael. "A two-armed bandit theory of market pricing." Journal of Economic Theory 9.2 (1974): 185-202.

Samuelson, Larry. "Information-Based Relative Consumption Effects." Econometrica 72.1 (2004): 93-118.

Spiegler, Ran. Bounded rationality and industrial organization. Oxford University Press, 2011.

Stevenson, Betsey, and Justin Wolfers. Economic growth and subjective well-being: Reassessing the Easterlin paradox. No. w14282. National Bureau of Economic Research, 2008.

Tulving, Endel, and Donald M. Thomson. "Encoding specificity and retrieval processes in episodic memory." Psychological review 80.5 (1973): 352.

Veenhoven, Ruut, et al. Happiness in nations: Subjective appreciation of life in 56 nations 1946–1992. Erasmus University Rotterdam, 1993.

Vigna, Stefano Della, and Ulrike Malmendier. "Paying not to go to the gym." The American Economic Review 96.3 (2006): 694-719.

Weiss, Yoram, and Chaim Fershtman. "Social status and economic performance:: A survey." European Economic Review 42.3 (1998): 801-820.