10.14754/CEU.2015.04

Essays on Consumer Search and Switching

by

András Péter Kiss

Sumbitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy at
Central European University

Supervisor: Péter Kondor

Budapest, Hungary

10.14754/CEU.2015.04

CENTRAL EUROPEAN UNIVERSITY
DEPARTMENT OF ECONOMICS

The undersigned hereby certify that they have read and recommend to the Department of Economics for acceptance a thesis entitled "Essays on Consumer Search and Switching" by András Kiss

Dated: March 30, 2015

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as dissertation for the degree of Doctor of Philosophy.

Chair of the Thesis Committee:______________________Julius Horvath

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as dissertation for the degree of Doctor of Philosophy.

Advisor:______________________ Peter Kondor

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as dissertation for the degree of Doctor of Philosophy.

Internal Examiner:______________________Sergey Lychagin

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as dissertation for the degree of Doctor of Philosophy.

External Examiner:______________________José Luis Moraga-González

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as dissertation for the degree of Doctor of Philosophy.

Internal Member:______________________Andrzej Baniak

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as dissertation for the degree of Doctor of Philosophy.

External Member:______________________László Paizs

10.14754/CEU.2015.04

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as dissertation for the degree of Doctor of Philosophy.

External Member: Szakadát László ______________________ László Szakadát

10.14754/CEU.2015.04

CENTRAL EUROPEAN UNIVERSITY
DEPARTMENT OF ECONOMICS

Author: András Péter Kiss
Title: Essays on Consumer Search and Switching
Degree: Ph.D.
Dated: April 24, 2015

Hereby I testify that this thesis contains no material accepted for any other degree in any other institution and that it contains no material previously written and/or published by another person except where appropriate acknowledgement is made.

Signature of the author: ....

András Péter Kiss

10.14754/CEU.2015.04

Abstract

The thesis consists of three single-authored chapters on consumer search and switching behavior. Chapter 1 looks at markets where information intermediaries, such as price comparison websites, help people in their search for the best offers. I build a theoretical model to examine the economic forces that drive the market structure of platforms towards, or away from, monopolization, and conclude that the latter ones will prevail.

Chapters 2 and 3 are empirical investigations on the switching behavior of consumers in the auto liability insurance market in Hungary. Both chapters are based on a unique, contract-level dataset that I collected from an insurance brokerage firm. In Chapter 2, I exploit a change in market regulation to estimate the causal effect of an advertising campaign on switching rates. The campaign's effect is large and mainly works through drawing people's attention to the switching opportunity.

In Chapter 3, I employ the dataset on insurance contracts to estimate switching costs in the market. I modify a standard two-period multinomial choice model by including the possibility of inattention to switching, which turns out to be a significant improvement to the econometric model's fit to the data and produces much more plausible results.

I provide more details on the contributions of the three chapters of the thesis below.

Chapter 1: Platform Competition with Price-Ordered Search

Chapter 1 contributes to the theoretical literature on competition between platforms in two-sided markets. I build a search model in which competing information platforms provide ordered price information to consumers, gathered from the sellers who voluntarily sign up for the advertising service. My central question is whether competition between platforms is sustainable, or the market will tip towards a monopoly provider. I find that multi-homing on the sellers' part, coupled with small frictions in buyer behavior, is sufficient to neutralize the positive cross-market externalities and prevent market tipping. When platforms can charge sellers for consumer traffic ("clicks"), the market achieves an efficient allocation. Without per-click charges, platforms will inefficiently restrict seller-to-seller competition to extract more surplus from consumers. The model's main application is the online comparison shopping industry.

Chapter 2: Salience and Switching

In Chapter 2, I estimate the effect of a consumer awareness campaign on contract switching decisions in auto liability insurance, and show that consumers' ignorance can be a major obstacle to switching service providers. For identification, I exploit a recent change in Hungarian regulation, which creates exogenous variation in the salience of the switching opportunity for a subset of drivers. Using a unique micro-level dataset collected from an insurance intermediary, I find that the campaign increases switching rates by 12 percentage points from a baseline of 20 percent. In comparison, the estimated reduced-form relationship between financial incentives

10.14754/CEU.2015.04

and switching decisions is much weaker: an additional saving of $50 per year - or about one-third of the median annual premium - is associated with only 4 percentage points higher switching rates. From a policy perspective, my results indicate that consumers could derive considerable benefits from effective information-spreading and market education campaigns, as well as a market design that makes infrequent, but economically significant choice situations more salient.

Chapter 3: Measuring Switching Costs in the Hungarian Auto Liability Insurance Market

Chapter 3 contributes to the literature on the measurement of consumer switching costs, using the dataset constructed in Chapter 2. Specifically, I estimate a structural model for contract switching in the Hungarian auto liability insurance market, allowing for switching costs and inattention to influence consumer decisions through separate channels. I find that inattention to the switching opportunity affects two-thirds of the population and explains consumer inertia to a large extent. Switching costs for attentive consumers are around $65, but are vastly overestimated when inattention is not accounted for. I also show that a concentrated media campaign can increase awareness to the switching opportunity by 23 percentage points from a baseline of 29 percent.

10.14754/CEU.2015.04

Acknowledgements

I am grateful to my supervisor, Péter Kondor, for his continuous support, valuable advice, and motivating presence throughout my many years of research.

Ádám Szeidl, Miklós Koren, and Botond Kőszegi have provided me with much needed guidance in different parts of the thesis. I am also indebted to my examiners, Sergey Lychagin and José Luis Moraga-González, for their useful comments and constructive criticism.

I am thankful to my professors and fellow students at the Economics Department for thoughtful discussions, and especially to dr. Gábor Vass and Árpád Antal for their support of the auto liability insurance project and their help with data acquisition. CEU has provided me with generous financial assistance during my studies, which made my life much easier.

Although I wrote every sentence of the dissertation myself, I can't help feeling that my parents also have a hand in it. I will always be grateful for the importance they placed on my education.

Mostly, I thank my wife for putting up with years of work that didn't quite seem like proper work, and my son for helping me keep a healthy perspective on all things throughout.

10.14754/CEU.2015.04

Contents

10.14754/CEU.2015.04

10.14754/CEU.2015.04

List of Tables

10.14754/CEU.2015.04

List of Figures

10.14754/CEU.2015.04

Chapter 1

Platform Competition with Price-Ordered Search

1.1 Introduction

New online technologies promised to eliminate the need to search all day long for the best price on a product. On comparison shopping sites, for example, all offers can be listed with lowest price first at a single click. The marginal cost of displaying another item to another consumer is negligible, which implies substantial economies of scale in providing comparative price information. Yet we do not observe *a* single price comparison website emerging; rather, a multitude of them fragment the market providing essentially the same service: showing ordered lists of offers to consumers.[1]

The lack of a single information "clearinghouse" seems all the more puzzling when we consider the network effects present in the market for price information. Since buyers prefer seeing more competing sellers and sellers value the visibility to more buyers, there should be a self-reinforcing tendency for a price comparison platform to accumulate clients on both sides, eventually tipping the market altogether towards a single monopoly provider.

I model the competition in online price comparison services and investigate the resulting equilibrium market structure. I find that the intuitively appealing positive network effects mechanism does not hold in this setting. Market concentration among comparison sites does not occur endogenously; in fact, the only market equilibrium is a symmetric one: websites will attract an equal number of buyers and sellers, and provide the same expected distribution of prices. This outcome holds with or without the ability of platforms to set traffic-related ("per-click") advertising fees and only requires the presence of slight informational frictions in consumers' choice of which platform to visit for price information. I also show that stronger frictions on the buyer side result in higher advertising fees for sellers, increase product prices and platform profits, and decrease consumer surplus. The overall welfare effect of frictions is negative when platforms cannot charge for traffic, but is neutralized by per-click pricing.

[1] In Germany, for example, there are at least a dozen general purpose price comparison sites. A number of these are listed at www.preisauskunft.de (3/4/2014).

10.14754/CEU.2015.04

In the model, buyers choose one of two price advertising platforms to search for the best offer on a product among the sellers who have signed up to the platform. Sellers can be present on more than one platform, but have to advertise the same prices in all locations. Buyers observe all prices on the chosen platform up front. Their surplus from a potential transaction also depends on a random binary match parameter (best thought of as a subjective assessment of seller trustworthiness). Buyers' platform choice is governed by a noisy signal of how much surplus they can expect to gain on either platform. Platforms earn money by charging sellers for the price listing service, but to do so successfully, they must also attract buyers with the promise of a range of good offers on the product.

The central insight of the model is that the ability of sellers to multi-home in two-sided markets provides a force against the self-reinforcing cross-market externalities, which is sufficient to prevent the endogenous formation of a monopoly platform. Non-exclusivity on the seller side enables platforms to copy each other's range of offers and hence be equally attractive to buyers.

In turn, small informational frictions in buyers' platform choice exclude equilibria in which all buyers and sellers coordinate to a single platform, because they hold the *ex ante* inexplicable belief that only this platform will succeed. Since platforms cannot "outdo" one another on the seller side, their competition for buyers reduces to a Cournot-type oligopoly game with a unique symmetric equilibrium and equal market shares. Frictions decrease the elasticity of demand for price comparison services and hence enable platforms to charge higher prices (to sellers) in equilibrium and indirectly extract more surplus from consumers.

A related non-tipping result arises from a modification of the baseline model, in which platforms can only levy listing fees, but are unable to charge sellers for consumer traffic. Here, platforms will be motivated to strongly restrict the number of sellers they show to consumers to limit seller-to-seller competition and hence extract more profit through the advertising fee. Total surplus will be decreased, because having fewer sellers implies a larger number of poor buyer-seller matches and unrealized transactions. At the same time, multi-homing loses some of its significance: when platforms only list a small fraction of available sellers, exclusivity is no longer a binding constraint of showing just as many offers as a rival site.

As an extension to the baseline model, I show that the non-tipping result can be overturned by assuming that sellers face fixed costs of platform participation, in addition to the fees charged by the platforms. Asymmetric equilibria exist in the model for a range of participation costs, because platforms that are larger on the buyer side allow sellers to re-capture their fixed costs with lower markups, which will, in turn, attract more buyers.

Applications of the presented platform competition model include websites built for everyday comparison shopping, such as shopping.com, pricerunner.co.uk, or billiger.de—and their numerous competitors—based in the U.S., U.K., and Germany, respectively. Each carry thousands of products in dozens of categories from electronics to household appliances, clothing, jewellery, and many others, and receive millions of unique visitors each month. Comparison sites specialized to more expensive, or one-off purchases, such as cars, airline tickets, or hotel rooms also abound, and generate similar traffic numbers. These sites have become an essential part of

10.14754/CEU.2015.04

the online consumer shopping scene and direct millions of non-trivial purchases on an everyday basis.

A common observation across all these examples is that there is no single dominant market player in information provision, that is, these platform markets do not "tip" to a monopoly in spite of the positive cross-platform externalities. In contrast, there are prominent examples for online marketplaces dominated by one platform: such as internet searches, consumer auctions, or social networks, to which my model does not apply.

This the first paper on clearinghouse-type search models that shows that several information clearinghouses can coexist in equilibrium without any dominant player endogenously driving the others out of business. Existing models in the literature are based on monopoly platforms by assumption, and tend to focus on other aspects of industry organization, such as price distortions on the two sides of the market (Baye and Morgan (2001)), product differentiation (Galeotti and Moraga-González (2009)), optimal advertising fee structure (Baye et al. (2011)), or price discrimination (Moraga-González and Wildenbeest (2012)). My model complements these papers by analyzing the competition of platforms in a similar search environment.

The structure of the chapter is the following. In section 2, I present the main ingredients of the baseline model with per-click fees. Section 3 contains the equilibrium solution, comparative statics, and welfare analysis. In section 4, I present two extensions. The first is an alternative pricing scheme, in which platforms are unable to charge for traffic. Although this assumption leads to the same conclusion of no market tipping, the mechanisms by which we arrive there are different. Second, I allow sellers to have fixed participation costs on platforms, which overturns the non-tipping results for certain parameter values. Section 5 contains a discussion of the chapter's results, followed by an exposition of related insights in the literature. Section 7 concludes. Longer proofs are left for the Appendix.

1.2 Model

Three groups of agents interact in the model: consumers (buyers), firms (sellers), and price comparison platforms. A close real-world analogy is that of intending to purchase a well-defined product, such as a specific TV set. Consumers might have already inspected the TV set at an electronics store in their neighborhood and have learnt its "High Street" price. They might also be aware that the product is sold online. However, discovering the available online offers without specialist guidance is costly, since online retailers are often small and hard to find.[2] For this reason, a consumer may turn to a—more visible—price comparison platform that provides an information channel between the two sides of the online market.

[2]Generic Internet search engines, for example, bring up all sorts of results related to search keywords, only some of which might correspond to real-time seller offers, the rest being product reviews, or other comparatively stale information.

10.14754/CEU.2015.04

1.2.1 Consumers

There are a continuum of buyers with a measure normalized to 1, each of whom is interested in buying a good that is uniformly valued at $r > 0$.[3] Every consumer visits a single price comparison site to search for offers of the product.

In addition to the price of the product, consumers also care about the shopping experience in general. In particular, they might be persuaded not to deal with a given seller if, on closer inspection, the details of the transaction process do not gain their trust.[4] The model captures this feature of consumer preferences by a random multiplicative "match value" that is specific to the buyer-seller pair, and is outside the influence of either party.

Specifically, when any consumer inspects any seller, there is a probability $0 < \beta < 1$ that the consumer will find the seller agreeable (yielding a match value multiplier of 1), in which case buying the product from the seller will result in a gross consumer surplus of r. With probability $1 - \beta$, the match value, and therefore also the gross surplus of the transaction, will equal 0. In the latter case, consumers will prefer not to purchase the product from the firm and will move on to inspect another seller.[5] If a consumer draws a zero match value for all sellers on the platform, he drops out of the online market and achieves zero net surplus in the model.[6]

For an example that generates such matching, consider the following. Let's assume that the platform operates an *ex post* evaluation system (as many platforms do), in which buyers leave feedback about their shopping experience with various sellers. All sellers are equally likely (with probability β) to satisfy an order without problems, in which case they get a positive evaluation. When a buyer inspects a seller, he looks at the last evaluation. If the evaluation is positive, the buyer goes through with the order. If the evaluation is negative, no order will be placed. Since the probability of reading a positive evaluation is β, the resulting matching process is equivalent to the model's assumptions.[7]

[3] This common valuation for the product can also be thought of as the universally known "High Street price", which is set to the demand of less sophisticated offline shoppers. I assume that this price is exogenous to the model, which simplifies the analysis, but also removes potentially interesting interactions between online and offline shoppers (see Baye and Morgan (2001) for more details).

[4] Certain consumers might, for example, dislike the payment, delivery, return, or warranty conditions of a given firm, which other consumers might find acceptable. People can also differ in their subjective assessments over the trustworthiness of sellers (deduced from the appearance and content of their webshop, for instance, or from a random sampling of the feedback provided by previous customers).

[5] Chen and He (2011) use a related assumption in the context of keyword-based paid-placement advertising. In their model, the match probability between an advertising firm and a random consumer is firm-specific and exponentially decreasing among advertising firms. In contrast, I assume that each seller is equally acceptable to a random buyer *ex ante*, and buyers' assessments are independent of each other after the inspection of a seller. Zhang (2009) also uses a similar demand structure in a sequential search model where consumers can observe prices before confirming the match quality.

[6] In reality, some people might continue searching for offers through other price comparison sites, but many would most likely give up hunting for bargains and simply visit the neighborhood store to buy the product at a price of r. With homogenous consumers, either everyone visits a single platform, or everyone searches on both platforms simultaneously. In the latter case, no seller would sign up to more than one platform, which contradicts the observation that multi-homing on the seller side is the rule rather than the exception in these markets. On balance, I found that the single-homing assumption on the buyer side is a more plausible description of online shopping.

[7] Although such *anecdotal reasoning* is not fully rational, there is ample psychological evidence that people are prone to this bias.

10.14754/CEU.2015.04

This type of matching describes one-off transactions better than recurring ones. In the latter case, one might imagine that buyers place a larger weight on their own previous experience with a seller than they do on others' experiences, which might lead to a different steady state than the one in the model. In this sense, the model applies to infrequent purchases of consumer electronics, for example, but not to frequently used food delivery services. Also, heterogeneity across sellers that is valued the same way by all buyers does not fit well with the matching process. For example, buyers probably all weakly prefer faster delivery to slower delivery for the same price, in which case their match value multipliers will not be independently drawn.

1.2.2 Firms

There are a non-trivial, but finite number ($N \gg 1$) of online firms selling a homogeneous product with identical and constant marginal cost $0 \leq m < r$. Firms might list prices on more than one platform (multi-home), but are not allowed to discriminate between platforms in their price setting.[8] If a seller abstains from price advertisements, it receives a profit normalized to 0 as an outside option. I assume that firms choose to advertise their prices on a platform if they expect non-negative profits from doing so.

1.2.3 Price comparison platforms

There are two price comparison platforms (indexed by j) which offer an informational channel through which sellers and buyers can connect in large numbers at a negligible cost per meeting. In return, platforms charge sellers a listing fee $a_j \geq 0$ and a per-click fee $c_j \geq 0$, both of which might differ across websites, but not across firms on the same website.[9] Transaction-based charges are not possible, since buyers purchase the items directly from the sellers, using the sellers' own webshops. However, the matching process results in constant clicks-to-sales conversion rates (equal to β), which leads to an equivalence between per-click charges and transaction fees in the model.

Charging consumers for viewing the list of prices is assumed to be impractical, which also coincides with what we observe in reality. Operating the price comparison site has no fixed cost (for simplicity), and any number of firms and consumers can be served at zero marginal cost. Platforms—just as firms—maximize their expected profits.

[8]This restriction captures the fact that price comparison platforms are normally information providers only, and the actual transactions take place in the sellers' webshops. Since buyers often leave and return to these webshops (even repeatedly) before purchasing an item, it is a real reputational hazard for a seller to confuse shoppers by advertising different prices on different platforms.

[9]Arguably, platforms could provide sign-up bonuses to firms, instead of charging them listing fees, leading to $a_j < 0$. But they would then be vulnerable to exploitation by any "seller" who chose to advertise a price above r, and hence never get (and pay for) any clicks, but could collect the sign-up bonus for free. To preclude such a possibility, I restrict the listing fee to be non-negative. A negative per-click charge could similarly be turned against a platform by the collusion of "sellers" and "buyers".

10.14754/CEU.2015.04

1.2.4 Shopping behavior

There are two aspects of the shopping process of consumers that need further specification: their choice of (1) which platform to visit, and (2) which seller to buy from on that platform. I look at each in turn.

Choice of platform

Rational buyers' platform choice should be governed by the net surplus they expect to gain on different platforms. However, the informational channel (online and offline advertising, news media, personal recommendations) through which they might learn about platforms contains a fair amount of noise. As a result, even if one price comparison site offers lower expected net surplus to buyers than the other, there is always a chance that a share of buyers will end up visiting the less valuable site. This share shrinks as the difference between platforms increases.

Formally, I assume that for the sake of the platform choice process, buyer b behaves *as if* he attached a random expected utility U_{bj} to platform j of the following form:

$$U_{bj} = u_j + \sigma \cdot \varepsilon_{bj} \tag{1.1}$$

where u_j is the "observable" part of utility and ε_{bj} is an independent random component. To arrive at closed-form solutions throughout the chapter, I assume that the random component follows a Type I extreme value distribution,[10] which turns the buyers' platform choice into a familiar multinomial logit problem.

The $\sigma > 0$ term is a scaling factor on the importance of the random component relative to the expected surplus provided by the comparison site. When σ is small, buyers put larger weight on the deterministic part of utility, u_j, which prompts them to make "sharper" choices between platforms based on the net surplus they expect to receive, and react more intensely to a change in this measure. The inverse of the parameter, $\frac{1}{\sigma}$, can therefore be interpreted loosely as the accuracy of the information that buyers possess about the platforms.

Although consumer behavior in the model is formally equivalent to a multinomial choice framework, I interpret ε_{bj} as an error in the choice process, and not as an unobserved preference parameter. Accordingly, u_j is included in the welfare measure, but $\sigma \cdot \varepsilon_{bj}$ is not.[11]

Since the measure of buyers is normalized to 1, the number of buyers visiting platform j (denoted by μ_j) is equal to the probability that platform j provides the highest utility to a given consumer. Following the derivation of choice probabilities in a multinomial logit model, the market share of platform 1 on the buyer side can be expressed as:

$$\mu_1 = \Pr\left(U_{b1} > U_{b2}\right) = \frac{\exp\left(\frac{u_1}{\sigma}\right)}{\exp\left(\frac{u_1}{\sigma}\right) + \exp\left(\frac{u_2}{\sigma}\right)}. \tag{1.2}$$

[10]That is, the CDF of ε_{bj} is given by $F^{EV}(x) = e^{-e^{-x}}$.

[11]Alternatively, I could assume that buyers have a preference for one platform over another for reasons other than the expected surplus from meeting sellers. An often-employed approach in the literature is to use a Hotelling-type model to capture preference heterogeneity (see Armstrong (2006), for example). However, as I stress the homogeneity of the price comparison service that platforms provide in the market, it would be counter-intuitive to assume that buyers have intrinsic preferences for platforms.

10.14754/CEU.2015.04

As the partial derivative

$$\frac{\partial \mu_1}{\partial u_1} = \frac{1}{\sigma} \mu_1 (1 - \mu_1) \tag{1.3}$$

shows, consumers react more sensitively in their website choice to a change in the expected net surplus when σ is lower. In this sense, σ is also a measure of "frictions" in consumer shopping behavior. As frictions disappear ($\sigma \to 0$), all informed consumers rush to the objectively more attractive website , whereas at the other extreme ($\sigma \to \infty$), all sites get an equal share of visitors irrespective of the expected surplus they provide.

Choice of transaction partner

The second aspect of consumers' shopping behavior concerns what they do on a price comparison platform. Since they can expect equally high *gross* surplus at any seller *ex ante*, it is rational for them to start inspecting the lowest pricing seller first (usually found on top of the offer list) and move down the pricing order until they find a suitable deal.[12] Charging a low price can therefore be interpreted as buying prominence in the consumer search process, a phenomenon investigated from a different angle by Armstrong et al. (2009) and Armstrong and Zhou (2011).

Consequently, the lowest pricing seller on site j can expect $\beta \cdot \mu_j$ actual buyers, the second down the list $(1 - \beta)\, \beta \cdot \mu_j$, the third $(1 - \beta)^2\, \beta \cdot \mu_j$, and so on. In the end, a fraction $(1 - \beta)^{n_j}$ of consumers will not find a suitable match on the site and will give up online shopping empty-handed (where n_j is the number of sellers on site j). Since $0 < \beta < 1$, this fraction shrinks as the number of advertising firms increases on a platform.[13]

1.3 Equilibrium

The timing of the game is as follows. First, comparison sites simultaneously set advertising (listing and per-click) fees, anticipating the reaction of firms and consumers. Second, firms simultaneously decide whether they wish to advertise their prices, and if yes, then pick one or both platforms for doing so. Without observing what sellers have actually done, each site is visited by a fraction of the consumers given by (1.2). As a result, sellers take both the distribution of consumers and of rival sellers across platforms as given when making their advertising and pricing decisions. Buyers inspect offers one-by-one, and when they find a suitable match they purchase the product.

I proceed to solve for the equilibrium of the model in two stages. First, I look at the optimal advertising and price setting response of sellers to the distribution of consumers and advertising fees across the two platforms. Second, I solve for the

[12] If two or more sellers are tied at the same price, I assume that they are investigated in a random order by the consumers who have not yet found a match.

[13] Ellison and Ellison (2009) estimate a demand model for memory chips sold through a price comparison site and find that the rank of a seller has a strong positive influence on the amount of transactions it conducts. In their estimates, being in the 7th position on the price list instead of the first place decreases sales by 83 percent, *ceteris paribus*. In my model, this finding would imply a match probability of $\beta \approx 0.25$.

optimal fee setting decision and the resulting market equilibrium of the platforms, taking into account the reaction of both sellers and buyers.

1.3.1 Advertising and pricing by sellers

Suppose that platform j provides access to μ_j customers, has a single seller, and charges a listing fee of a_j and a per-click fee of c_j. The expected profit of the seller charging a price p for the product is then given by:

$$\pi_j(p) = (p - m)\,\beta\mu_j - c_j\mu_j - a_j \tag{1.4}$$

This expression is clearly maximized at $p = r$: a single seller will always set a monopoly price.

For this setup to be an equilibrium in seller strategies, it must also be the case that $\pi_j(r) = 0$, otherwise at least one more seller would want to join the platform and slightly undercut the first entrant. This zero-profit condition yields a profit of:

$$\Pi_j = c_j\mu_j + a_j = (r - m)\,\beta\mu_j \tag{1.5}$$

for the platform and expected consumer surplus of $u_j = 0$.

As expression (1.4) shows, there are two ways in which a single seller who plays a pure pricing strategy can earn zero profits on a comparison site. The site can charge a maximal per-click fee of $c_j = \beta(r - m)$ and a listing fee of $a_j = 0$, or a per-click fee of $c_j \in [0, \beta(r - m))$ and a listing fee of $a_j = [\beta(r - m) - c_j]\,\mu_j > 0$. The two cases have different implications for the entry of rival sellers and the resulting equilibrium.

The first case of charging maximal per-click fees leads to an effective marginal cost of $m + \frac{c_j}{\beta} = r$,[14] which leaves no room for a seller to price the product below r, because lower prices would always yield negative profits. As a result, the entry of additional sellers to the platform is not hindered by the prospect of Bertrand competition. If a second seller were to enter, it would also set a price of r, earn a zero expected margin on each sale, but since the listing fee is also zero, the new entrant would not end up with negative profits. Hence, by the sellers' entry rule, all N sellers would want to advertise their—monopoly—prices on the site, yielding a profit of:

$$\Pi_j = \left[1 - (1 - \beta)^N\right](r - m)\,\mu_j$$

for the platform and expected consumer surplus of $u_j = 0$. $(1 - \beta)^N$ is the probability that a buyer finds no match among N sellers, hence the total surplus created by transactions on the platform equals $\left[1 - (1 - \beta)^N\right](r - m)\,\mu_j$, all of which the platform appropriates through advertising fees. With such an outcome, the platform only provides a seller location service for buyers, but prevents seller-to-seller competition altogether.

In what follows, we will see that this maximal per-click fee outcome also arises as a degenerate equilibrium in a more general game with mixed pricing strategies on the sellers' part.

[14]Since only a β fraction of visits are turned into sales, a seller must pay the platform the price of $\frac{1}{\beta}$ "clicks" for each sale, on average.

In the second case of less-than-maximal per-click fees, entering sellers would always have an incentive to either slightly undercut one another, or to raise prices and benefit from sales to consumers who were poorly matched with low pricing firms. So even if the comparison site chose to reduce its listing fee to zero to attract all sellers, no equilibria could exist in pure pricing strategies. There might, however, still be mixed strategy equilibria, in which sellers randomize over prices. This result is summarized in the following proposition.

Proposition 1.1. *With per-click charges of $c_j < \beta(r-m)$ and any feasible listing fee in the interval $[0, \beta(r-m)\mu_j - c_j\mu_j]$, no pure strategy equilibrium exists for advertising on platform j. If we allow sellers to employ symmetric mixed strategies in price setting, then the only existing equilibrium must involve sellers choosing prices from a continuous price distribution that contains no mass points and has a support on $\left[\underline{p}_j, r\right]$ where $m + \frac{c_j}{\beta} < \underline{p}_j < r$. In particular, the reservation price of consumers is always charged with a positive density in a symmetric mixed strategy equilibrium.*

Proof. For the detailed argument, which is analogous to the one in Varian (1980) and Baye and Morgan (2001), see the set of lemmas in the Appendix. □

It is an important feature of the pricing game that advertising sellers charge the reservation price of buyers with a positive density. The reason is the following. When the number of rival sellers is finite, a seller charging the highest price in the (symmetric) equilibrium strategy will end up in the last position of the comparison site's price list with probability 1 and its expected sales will be independent of its price. Therefore, its optimal action is to set this "last-in-line-for-sure" price equal to the buyers' reservation value r.

Characterizing the symmetric mixed strategy equilibrium

Proposition 1.1 established that if $c_j \leq \beta(r-m)$, we should be looking for an equilibrium price distribution F_j that all $n_j \geq 2$ firms play simultaneously on site j.[15] The next proposition characterizes the equilibrium behavior of sellers on a platform for given advertising charges.

Proposition 1.2. *When a platform has μ_j buyers, sets per click charges of $c_j \leq \beta(r-m)$ and listing fees of a_j, the number of sellers on the platform (n_j) is determined by the advertising demand function*

$$(1-\beta)^{n_j-1}\beta\mu_j\left(r-m-\frac{c_j}{\beta}\right) = a_j \tag{1.6}$$

[15] For ease of exposition, I will derive the equilibrium assuming that sellers advertise on platform j only. In a symmetric equilibrium with multi-homing (the only equilibrium of the model), the price distributions are identical to the single-homing case.

Moreover, sellers play a mixed pricing strategy according to the cumulative distribution function

$$F_j(p) = \frac{1}{\beta} - \frac{1-\beta}{\beta}\left(\frac{r-m-\frac{c_j}{\beta}}{p-m-\frac{c_j}{\beta}}\right)^{\frac{1}{n_j-1}} \tag{1.7}$$

which has full support on $\left[\underline{p}_j, r\right]$, *where*

$$\underline{p}_j = \left(m+\frac{c_j}{\beta}\right) + (1-\beta)^{n_j-1}\left[r-\left(m+\frac{c_j}{\beta}\right)\right] \tag{1.8}$$

Proof. Let us first look at what happens when a firm charges the highest price r, from the support of F_j. Since F_j is atomless, the high-pricing seller will be at the end of the price list with probability 1, and will receive $(1-\beta)^{n_j-1}\cdot\mu_j$ clicks, β share of which will turn into sales with margin $r-m$. The seller's expected profit is therefore:

$$\pi_j(r) = \beta(1-\beta)^{n_j-1}\mu_j\cdot(r-m) - (1-\beta)^{n_j-1}\mu_j\cdot c_j - a_j \tag{1.9}$$

In equilibrium, this profit must be equal to zero (the firms' outside option), which yields the implicit expression in equation (1.6) for the demand for advertising on platform j. As expected of a demand function, the number of advertising firms (n_j) varies inversely with both the listing (a_j) and the per-click (c_j) charges. Since the right hand side of (1.6) is non-negative, the per-click fee cannot exceed $\beta(r-m)$ for any seller to be interested in listing prices on site j.

Expected seller profits must also equal zero for all other prices that are played with positive density in a mixed strategy equilibrium, which helps us pin down the entire price distribution F_j as follows.

The expected profit of a seller when it advertises a price p on site j can be expressed as:

$$\pi_j(p) = \sum_{k=1}^{n_j}\left\{\binom{n_j-1}{k-1} F_j(p)^{k-1}\left[1-F_j(p)\right]^{n_j-1-(k-1)}\right\} \cdot\left\{(1-\beta)^{k-1}\beta\mu_j\right\}\left(p-m-\frac{c_j}{\beta}\right) - a_j \tag{1.10}$$

The expression within the first set of braces is the probability that the seller ends up in the kth place on the list with n_j-1 rivals. The second set of braces contains the expected number of consumers buying from a seller who is in the kth place, followed by the per sale margin $\left(p-m-\frac{c_j}{\beta}\right)$, and the listing fee a_j.

Using the Binomial Theorem, (1.10) can be simplified to:

$$\pi_j(p) = \left[1-\beta F_j(p)\right]^{n_j-1}\beta\mu_j\left(p-m-\frac{c_j}{\beta}\right) - a_j. \tag{1.11}$$

10.14754/CEU.2015.04

In particular, since $p = r$ was shown to be included in the support of F_j and $F_j(r) = 1$, we have:

$$\pi_j(r) = (1-\beta)^{n_j - 1}\beta\mu_j\left(r - m - \frac{c_j}{\beta}\right) - a_j$$

just as in (1.9). Setting $\pi_j(p) = \pi_j(r)$ for all prices in the support of F_j, we arrive at the equilibrium cumulative price distribution function in (1.7), along with the associated probability density function:

$$f_j(p) = \frac{1}{n_j - 1}\frac{1-\beta}{\beta}\left(r - m - \frac{c_j}{\beta}\right)^{\frac{1}{n_j-1}}\left(p - m - \frac{c_j}{\beta}\right)^{-\frac{n_j}{n_j-1}} \tag{1.12}$$

and the lowest price charged with positive density in equilibrium, shown in (1.8). □

In particular, (1.8) shows that the lower end of the price distribution will get close to the effective marginal cost of the product $\left(m + \frac{c_j}{\beta}\right)$ as the number of competing sellers increases on a platform, but will not reach it as long as n_j remains finite. At the other end of the spectrum, the price distribution collapses onto the monopoly price $\left(\underline{p}_j = r\right)$ when there is only a single seller, or when the per-click fee approaches its maximal value of $\beta(r - m)$.

Single-homing versus multi-homing

The above analysis of seller strategies assumes that sellers only post prices on at most one of the two platforms, that is, they single-home. In this section, I will extend the analysis to multi-homing by looking at possible advertising and pricing decisions on both platforms.

In the pure strategy equilibrium of maximal per-click charges ($c_1 = c_2 = \beta(r - m)$) and zero listing fees, it doesn't matter whether the same seller is present on both sites, or different sellers post prices on different sites. Since all listed prices equal r on both platforms, the restriction of no cross-platform price discrimination does not disadvantage a multi-homing firm relative to two single-homing firms. Sellers are assumed to enter a comparison site when it promises non-negative profits, so all sellers will enter both sites.

The analysis gets only slightly more involved with mixed pricing strategies. Let us assume that platforms 1 and 2 charge advertising fees of (c_1, a_1) and (c_2, a_2), have customer bases of μ_1 and μ_2, single-homing sellers of n_1^S and n_2^S, and multi-homing sellers numbering n^M.[16] In equilibrium, when a single-homing seller charges the monopoly price r on platform j, it must earn zero expected profits:

$$\pi_j^S(r) = (1-\beta)^{n^M + n_j^S - 1}\beta\mu_j\left(r - m - \frac{c_j}{\beta}\right) - a_j = 0$$

[16]Platform j thus has $n_j = n^M + n_j^S$ sellers in total.

and the same must be true for a multi-homing seller concerning the sum of profits on both sites:

$$\pi^M(r) = \sum_{j=1}^{2} \pi_j^S(r) = 0$$

Therefore, the restriction on multi-homing firms of no cross-platform price discrimination makes the entry of multi-homing sellers neither harder, nor easier, than single-homing ones. That is, a multi-homing firm earns the same profit as two single-homing firms: zero. Consequently, the relative share of single- and multi-homing firms is undetermined in a mixed strategy equilibrium.[17] I will therefore not distinguish between single- and multi-homing sellers in subsequent derivations.

The price discrimination restriction might make a difference for the equilibrium single- and multi-homing price distributions that form the basis of randomization over product prices. However, the distributions themselves arise a consequence of the zero profit condition on firms, and not as a determinant of profits. Hence no seller would be concerned about differences between two mixed strategies, as long as both yield the same profit.[18]

1.3.2 Maximizing platform profits

Propositions 1.1 and 1.2 established the equilibrium behavior of sellers on a comparison site, given the advertising charges set by the site and the distribution of consumers across sites. In this section, I will close the model by characterizing the optimal price setting strategies of platforms and the resulting equilibrium. The next proposition summarizes the results.

Proposition 1.3. *The model has a unique equilibrium, which is symmetric in the platforms' actions. When consumers are sufficiently uninformed about platforms, such that σ is below a threshold value $\overline{\sigma}_l$, per-click fees are maximal and consumers receive zero expected utility in equilibrium. When they are sufficiently well-informed, such that $\sigma > \overline{\sigma}_u$, per-click fees are zero and consumer utility is maximized given the total number of sellers. In the intermediate range of $\overline{\sigma}_l \leq \sigma \leq \overline{\sigma}_u$, equilibrium consumer utility varies inversely with the informational friction parameter σ.*[19]

The proof of the proposition is separated into a series of lemmas below.

Lemma 1.1. *Platform profits can equivalently be written as function of advertising charges, or as a function of the number of sellers listed and expected utility provided to consumers.*

Proof. The expected profit of price comparison platform j is given by:

$$\Pi_j(c_j, a_j) = n_j \cdot a_j + \mu_j \gamma_j c_j \tag{1.13}$$

[17] Unless, of course, both platforms want to list all sellers in equilibrium, in which case each seller will choose to multi-home.

[18] Nevertheless, in a symmetric equilibrium all sellers employ identical distributions, regardless of single- or multi-homing.

[19] The specific values for the thresholds are $\overline{\sigma}_l = \frac{1}{2}(1-\beta)^{N-1} N\beta(r-m)$ and $\overline{\sigma}_u = \frac{1}{2}\left[1-(1-\beta)^N\right](r-m)$.

which is the objective function to be maximized by site j with respect to a_j and c_j, given the constraints imposed by the advertising demand in (1.6) and buyers' choice between comparison sites (see (1.2)). γ_j denotes the expected number of "clicks" that a buyer makes on the site, which is equal to:

$$\gamma_j = 1 + (1-\beta) + (1-\beta)^2 + \cdots + (1-\beta)^{n_j - 1} = \frac{1-(1-\beta)^{n_j}}{\beta} \tag{1.14}$$

Instead of directly maximizing (1.13) with respect to the listing and per-click fees, let us express the comparison site's profit in terms of the number of its sellers and the expected utility it provides to a consumer.

First, by substituting the advertising demand function in (1.6) and the expected number of "clicks" in (1.14) into the objective function, we get:

$$\Pi_j(c_j, n_j) = n_j (1-\beta)^{n_j - 1} \beta \mu_j \left(r - m - \frac{c_j}{\beta} \right) + \mu_j \frac{1-(1-\beta)^{n_j}}{\beta} c_j \tag{1.15}$$

Second, the expected surplus that a platform's entire consumer base receives equals the difference between the total expected surplus created by all transactions through the platform and the profit (i.e. total advertising revenue) of the comparison site:[20]

$$\mu_j \cdot u_j = \mu_j \left[1 - (1-\beta)^{n_j}\right](r-m) - \left[n_j \cdot a_j + \mu_j \gamma_j c_j\right] \tag{1.16}$$

where the first set of brackets contains the probability that a consumer finds a suitable match on the site, and the second set of brackets shows site j's profit. Substitution from (1.6) and (1.14) yields, after rearrangement and simplification, the following relationship between the expected utility of a single consumer and the decision variables n_j and c_j:

$$u_j = \left\{ 1 - (1-\beta)^{n_j} \left[1 + n_j \frac{\beta}{1-\beta} \right] \right\} \left(r - m - \frac{c_j}{\beta} \right) \tag{1.17}$$

For $n_j \geq 1$, u_j is positive and increasing in the number of sellers, provided that $c_j < \beta(r-m)$. At $c_j = \beta(r-m)$ and $a_j = 0$, sellers are indifferent between entering the platform and staying out, and when they do enter, they will charge the monopoly price, and hence leave zero surplus for consumers. For $c_j > \beta(r-m)$, no firm is prepared to advertise on platform j, again leading to zero consumer utility. Consumers therefore cannot do worse than $u_j = 0$, which is one boundary condition in the platform's profit maximization problem.

There is also a second, somewhat more subtle boundary condition on utilities. Since per-click charges cannot be negative, the maximum expected surplus a consumer can achieve on a platform is $\left\{ 1 - (1-\beta)^N \left[1 + N\frac{\beta}{1-\beta} \right] \right\} (r-m)$. The total

[20]The zero outside option of the sellers ensures that they do not keep any part of the total surplus in equilibrium.

expected surplus of a single consumer's transaction equals $\left[1-(1-\beta)^N\right](r-m)$, which is a strictly higher value than the upper bound on utilities. As a result, platforms always earn some profit in an equilibrium where the upper boundary condition is binding.

In the third step of transforming a platform's profit maximization problem, we can express the per-click fee c_j from (1.17) and plug it into the objective function (1.15) to get a relatively simple maximand:

$$\max_{n_j,u_j} \Pi_j = \mu_j \left\{\left[1-(1-\beta)^{n_j}\right](r-m) - u_j\right\} \tag{1.18}$$

subject to:

$$0 \leq u_j \leq \left\{1-(1-\beta)^N\left[1+N\frac{\beta}{1-\beta}\right]\right\}(r-m) \tag{1.19}$$

and also satisfying the consumers' platform choice rule given by (1.2). □

In what follows, I first derive the internal solution to the model, then describe what happens at the upper and lower utility bounds.

Internal solution

Lemma 1.2. *Platform profits are increasing in the number of sellers listed, hence both platforms will want to list all sellers.*

Proof. The first derivative of the objective function in (1.18) with respect to the number of sellers n_j is given by

$$\frac{\partial \Pi_j}{\partial n_j} = -\log(1-\beta)\,\mu_j\,(1-\beta)^{n_j}(r-m) \tag{1.20}$$

Since $0<\beta<1$, this expression is always positive, therefore the profit maximizing number of sellers equals N. □

As a result of this lemma, platforms have a dominant strategy of setting listing charges low enough to attract all sellers. They will therefore use their other instrument, the per-click fee, to compete for buyers. By equation (1.17), we can equivalently say that the competition for the buyer side is driven by the amount of surplus that consumers receive. The following three lemmas state the result of this rivalry.

Lemma 1.3. *Within the bounds given by* (1.19)*, the best response of platforms in terms of expected utility provided to consumers is increasing in each other's actions.*

Proof. Setting the partial derivative of (1.18) with respect to consumer utility equal to zero for profit maximization:

$$\frac{\partial \Pi_j}{\partial u_j} = \frac{\partial \mu_j}{\partial u_j}\left\{\left[1-(1-\beta)^{n_j}\right](r-m) - u_j\right\} - \mu_j = 0 \tag{1.21}$$

Using (1.2) and (1.3), we arrive at platform j's reaction function to platform i's utility choice:

$$\frac{1}{\sigma}\frac{1}{\exp\left(\frac{u_j-u_i}{\sigma}\right)+1}\left\{\left[1-(1-\beta)^N\right](r-m)-u_j\right\}=1 \quad i,j\in\{1,2\}, i\neq j \tag{1.22}$$

The left hand side is strictly decreasing in u_j and increasing in u_i, the reaction function is therefore positively sloped. □

Lemma 1.4. *The best response functions in* (1.22) *have a unique equilibrium, which is symmetric in the expected utility provided to consumers by the platforms and given by*

$$u_1=u_2=\left[1-(1-\beta)^N\right](r-m)-2\sigma \tag{1.23}$$

provided that $\overline{\sigma}_l \equiv \frac{1}{2}(1-\beta)^{N-1}N\beta(r-m)\leq\sigma\leq\frac{1}{2}\left[1-(1-\beta)^N\right](r-m)\equiv\overline{\sigma}_u$. *As a result, the market is always shared between comparison sites equally on the buyers' side.*

Proof. Examining (1.22), we can make the following observations. As $u_i\to\infty$, $u_j\to\left[1-(1-\beta)^N\right](r-m)-\sigma$, which may, or may not, be within the upper and lower bounds on u_j, as given by constraints in (1.19). The slope of the reaction function, by total differentiation:

$$\frac{\mathrm{d}u_j}{\mathrm{d}u_i}=\frac{\exp\left(\frac{u_j-u_i}{\sigma}\right)}{\exp\left(\frac{u_j-u_i}{\sigma}\right)+1}<1 \tag{1.24}$$

Examining both (1.22) and (1.24), we can see that the slope is decreasing as we move to the right on platform j's reaction curve (i.e. as u_j is increasing). Since the reaction function's slope is always less than 1, the number of times it can cross the $u_j=u_i$ (45-degree) line is either zero or one. We get one intersection if $u_j(u_i=0)>0$, and zero otherwise.

Combining the reaction function in (1.22) with the non-negativity constraint on utilities, it follows that if the parameters in (1.22) are such that $u_j(u_i)=0$ for some $u_i=\overline{u}_i\geq 0$, then for all $0\leq u_i\leq\overline{u}_i$, $u_j(u_i)=0$. In this case, $u_j=u_i=0$ are mutual best responses.

By the symmetry of the reaction functions, we have therefore shown both the uniqueness and the symmetry of the equilibrium. Algebraically, we arrive at (1.23) by setting $u_i=u_j$ in (1.22) and solving the equation for u_j. Combining (1.23) with the bounds on utilities in (1.19), we can derive the lower and upper thresholds $\overline{\sigma}_l$ and $\overline{\sigma}_u$ for the internal solution to apply. □

The symmetric reaction functions in (1.22) are shown graphically in Figure 1.1 for chosen parameter values. The total surplus of a transaction is $r-m=1$. On the left-hand side, consumers are relatively well-informed about platforms (σ is

10.14754/CEU.2015.04

Figure 1.1: Platform reaction functions in consumer utility space ($r = 1$, $m = 0$, $N = 100$, $\beta = 0.5$)

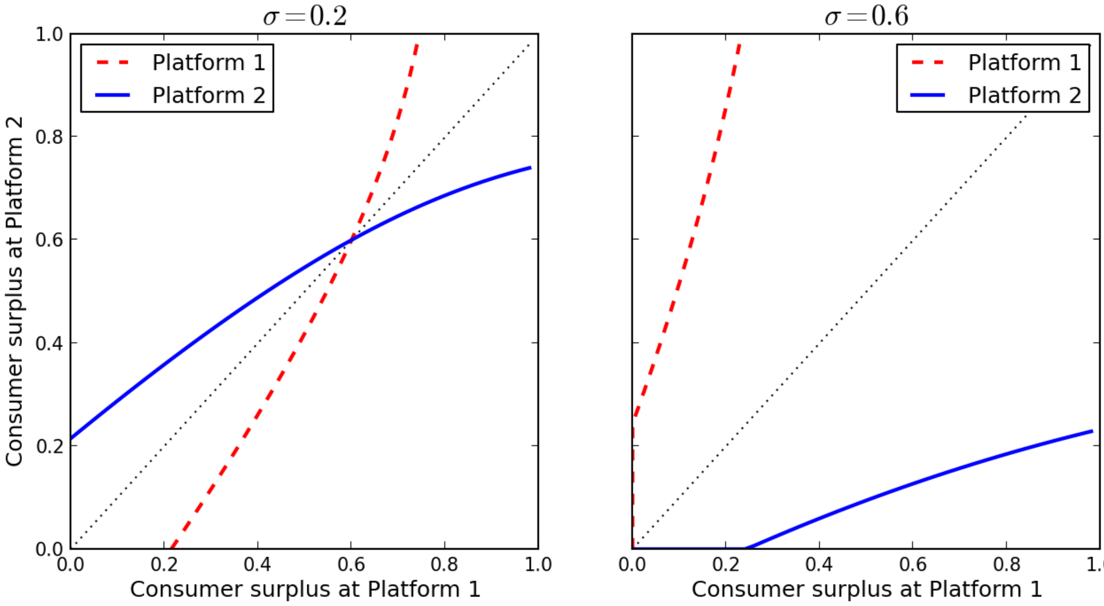

low), and therefore receive a large part of the total surplus in equilibrium. On the right hand side, consumers choose platforms with more noise, and end up with zero expected utility.[21]

Platforms leave more of the total surplus to consumers when consumers are better informed about their choices (σ is small). Total surplus itself is larger when there are more sellers (N is higher) and buyers are less picky about sellers (β is higher).

As the number of sellers becomes large, the lower boundary $\overline{\sigma}_l$ converges to zero and effectively disappears (for we have already constrained σ to be positive). The upper threshold, on the other hand, never exceeds $\frac{r-m}{2}$, and might therefore be binding with any number of sellers.[22]

Since the comparison sites will try to sign up all potential sellers ($n_1 = n_2 = N$), we can already solve for the equilibrium per-click charges using (1.17):

$$c_1 = c_2 = \beta \frac{2\sigma - (1-\beta)^{N-1} N\beta (r-m)}{1 - (1-\beta)^N \left[1 + N\frac{\beta}{1-\beta}\right]} \tag{1.25}$$

Finally, listing fees are set by the sites to motivate the entry of all sellers by allowing them nonnegative (in effect: zero) expected profits. Substitution from (1.25) into (1.6) yields the following expression:

$$a_1 = a_2 = \frac{\beta}{2}(1-\beta)^{N-1} \frac{\left[1 - (1-\beta)^N\right](r-m) - 2\sigma}{1 - (1-\beta)^N \left[1 + N\frac{\beta}{1-\beta}\right]} \tag{1.26}$$

[21] If σ exceeded 1 by a small amount, the reaction functions would fully coincide with the horizontal and the vertical axes, meaning that platforms would have a dominant strategy of not leaving any surplus for the consumers.

[22] In terms of the parameter values in Figure 1.1, an internal solution exists for $0 \lessapprox \sigma \lessapprox 0.5$.

10.14754/CEU.2015.04

Corner solutions

Lemma 1.5. *When consumers are extremely well, or extremely badly, informed, such that either* $\sigma < \frac{1}{2}(1-\beta)^{N-1} N\beta (r-m)$ *or* $\sigma > \frac{1}{2}\left[1-(1-\beta)^N\right](r-m)$*, two kinds of corner solutions can arise: a competitive and a non-competitive one. The first is characterized by zero per-click fees and consumer surplus being at its upper bound given by* (1.19). *In the non-competitive corner solution, per-click fees are maximal, expected consumer surplus is zero, and platforms share all the surplus of the game between themselves.*

Proof. The claims follow from the substitution of the utility bounds in (1.19) into the equilibrium expression given by (1.23).

In the first case, competition for consumers between comparison sites is intense. Platforms will decrease their per-click fees to zero, and would even prefer to make them negative, if exploitation by colluding sellers and buyers was not an issue.[23] However, sellers will still be charged positive listing fees of

$$a_1 = a_2 = \frac{\beta}{2}(1-\beta)^{N-1}(r-m)$$

to extract their profits, and therefore platforms also earn positive profits, even as $\sigma \to 0$.

In the second case, comparison sites are best off not competing for buyers by leaving them zero expected utility. Per-click fees will therefore be maximal:

$$c_1 = c_2 = \beta (r-m),$$

all sellers will charge the monopoly price, and pay nothing for appearing on the platforms ($a_1 = a_2 = 0$). This non-competitive equilibrium corresponds to the only existing pure strategy pricing outcome described before. □

Lemmas 1.1-1.5 also complete the proof of Proposition 1.3, providing the main result of the paper.

1.3.3 Comparative statics

Internal solution

The equilibrium per-click fee in (1.25) is decreasing in the precision of consumers' information about platforms, whereas the listing fee moves in the other direction. Stronger competition for buyers drives down the per-click fee, allowing the listing fee to rise without violating the participation constraint of sellers.

In general, the partial effects of the buyer-seller match parameter (β) and the number of sellers (N) on the platforms' fee structure are ambiguous. However,

[23]Instead, in the real world platforms will likely keep competing for consumers by enhancing their search experience, for example.

10.14754/CEU.2015.04

if there are enough sellers to neglect the possibility of not finding a match on a platform,[24] we are left with two simple expressions:

$$c_1 = c_2 \approx 2\beta\sigma \tag{1.27}$$

and $a_1 = a_2 \approx 0$. The per-click fees in the $N \to \infty$ limit are proportional to both the buyer-seller match parameter and the uninformedness of buyers as measured by the standard deviation parameter σ.

Intuitively, with a large number of sellers, consumers will be able to buy the item on platform j at a price very close to the effective marginal cost of $m + \frac{c_j}{\beta}$. A platform therefore competes by setting this marginal cost lower than its rival, but not quite in a pure Bertrand fashion, since the lack of perfect buyer information means that small decreases are not always noticed (and larger ones are not always profitable). The quality of buyer information (σ) thus determines the equilibrium effective marginal cost. Keeping this cost constant requires per-click fees that are proportional to the buyer-seller match parameter β.

Corner solutions

At a corner solution, per-click fees are either minimal (with well-informed consumers) or maximal (with badly-informed consumers), but not otherwise dependent on the information parameter σ. In the latter case, per-click fees are again proportional to the match parameter β, such that the effective marginal cost of the product remains constant (and equal to the reservation price r). Listing fees always adjust to extract all profit from sellers.

1.3.4 Welfare

In the model, total surplus equals the expected number of buyer-seller transactions multiplied by the net surplus of a single transaction:

$$W = \left[1 - (1-\beta)^N\right] \cdot (r - m)$$

Since sellers end up getting their outside option (zero profit), all surplus is shared between the consumers (S) and the two platforms (Π_1, Π_2): $W = S + \Pi_1 + \Pi_2$. The division rule is determined by the quality of buyer information about platforms in a non-monotonic fashion, as described by the following proposition.

Proposition 1.4. *At the internal solution characterized by* $\frac{1}{2}(1-\beta)^{N-1} N\beta (r-m) \leq \sigma \leq \frac{1}{2}\left[1-(1-\beta)^N\right](r-m)$, *each comparison site earns a profit equal to* σ, *and buyers keep the rest:*

$$\Pi_1 = \Pi_2 = \sigma$$
$$S = \left[1-(1-\beta)^N\right] \cdot (r-m) - 2\sigma$$

[24] Suppose that $\beta = 0.25$, a seemingly realistic value according to the analogous estimates of Ellison and Ellison (2009). For $N = 17$ sellers (taken as the average from a large sample analyzed by Baye et al. (2004)), the possibility of not finding a match on a platform is $(1-0.25)^{17} = 0.0075$. Therefore, the approximation is likely to be valid in reality.

When consumers choose platforms with too much noise (corner solution with a high σ), platforms divide all the surplus between themselves:

$$S = 0, \;\; \Pi_1 = \Pi_2 = \frac{1}{2}\left[1 - (1-\beta)^N\right] \cdot (r - m)$$

At the other extreme, if buyers' information is precise enough, platforms still receive a minimum profit level:

$$S = \left\{1 - (1-\beta)^N \left[1 + N\frac{\beta}{1-\beta}\right]\right\}(r-m)\,, \;\; \Pi_1 = \Pi_2 = \frac{1}{2}(1-\beta)^{N-1} N\beta\,(r-m)$$

Proof. Substitute the unique solution of the game given in Lemmas 1.2 and 1.4 into the expression for platform profit in (1.18). □

Buyers only benefit from the presence of more sellers when they are sufficiently well-informed about platforms. The marginal surplus of an additional seller is fully captured by buyers in the internal solution and by platforms in the non-competitive corner solution. Similarly, a welfare-enhancing increase in the buyer-seller match parameter β benefits only buyers in the internal solution and only platforms in the non-competitive corner solution.

The competitive corner solution is more complex in this respect, since the presence of additional sellers, or an increase in the match probability, initially benefit, but then harm platforms, whereas both effects are beneficial for consumers (for $N \geq 2$). For the comparison sites, the profit maximizing parameter values are:

$$N = \frac{1}{|\log(1-\beta)|}$$

and

$$\beta = \frac{1}{N}.$$

For $\beta = 0.25$, for example, platforms' profits start decreasing if there are more than 3 sellers in total.[25] Conversely, with 10 sellers, platforms would prefer a buyer-seller match probability of $\beta = 0.1$.

1.4 Extensions

1.4.1 No per-click charges

The baseline model with two-part tariffs yields the stark result that both comparison sites will want to sign up all sellers in the market by charging low enough listing fees to induce sellers to multi-home. On the buyer side, platforms compete in utilities offered to consumers, the equilibrium value of which depends on the elasticity of consumer demand, i.e. on how well-informed consumers are about the expected surplus on each platform. When platform competition is intense, per-click fees will

[25] Of course, this does not mean that competing platforms will only list 3 sellers when more are available. On the other hand, a monopoly platform would restrict the length of its price list to $\frac{1}{|\log(1-\beta)|}$.

drop to zero, although platforms will still extract some profit from sellers through the listing fee.

As we have seen, in the competitive corner solution platforms could benefit from having fewer sellers in the market. Unfortunately, decreasing their own seller base (by increasing the listing fee) won't do, because that would cause a drop in consumer utility, which is more profitably achieved by increasing per-click charges instead (but in a competitive corner solution, per-click fees must be zero).

This seller-base-reducing motivation of comparison sites in the competitive corner solution suggests that we might get markedly different outcomes if per-click pricing was not available at all.[26] In the current section, I analyze such a scenario.

If $c_j = 0$ by construction, advertising demand and expected consumer utility simplify to the following expressions (see (1.6) and (1.17)):

$$(1-\beta)^{n_j-1}\beta\mu_j(r-m) = a_j \tag{1.28}$$

and

$$u_j = \left\{1-(1-\beta)^{n_j}\left[1+n_j\frac{\beta}{1-\beta}\right]\right\}(r-m) \tag{1.29}$$

Comparison sites only earn money from listing fees, which modifies their objective function to:

$$\max_{n_j}\Pi_j = n_j a_j = n_j(1-\beta)^{n_j-1}\beta\mu_j(r-m) \tag{1.30}$$

subject to the consumers' platform choice behavior in (1.2), which is also indirectly influenced by n_j through consumer utility (see (1.29)). Differentiating with respect to the only decision variable n_j, we arrive at the comparison sites' first order condition:

$$\underbrace{(1-\beta)^{n_j-1}\beta\mu_j(r-m)}_{\text{profit extracted from an additional seller}} + \underbrace{[\log(1-\beta)]\cdot n_j(1-\beta)^{n_j-1}\beta\mu_j(r-m)}_{\text{more rivalry decreases total seller profits}} + \underbrace{n_j(1-\beta)^{n_j-1}\beta\frac{\partial\mu_j}{\partial u_j}\frac{\partial u_j}{\partial n_j}(r-m)}_{\text{new buyers increase total seller profits}} = 0 \tag{1.31}$$

The three components on the left hand side of (1.31) have straightforward interpretations. The first element is the expected transaction profit per seller, which is also equal to the listing fee that an additional seller would pay to the platform. The second component denotes the decrease in total seller profits (and therefore platform revenues) as a result of increased competition generated by the additional seller ($\log(1-\beta) < 0$), while the last component shows the increase in total seller profits following the increase in the number of buyers visiting the platform.

Equation (1.31) can be simplified by dividing both sides by the strictly positive expression for total platform revenues $R_j = n_j \cdot (1-\beta)^{n_j-1}\beta\mu_j(r-m)$. In addition, the proportional change in visitor numbers can be more specifically expressed using the consumer reaction function enbodied in (1.2) and (1.3), yielding the first order

[26] As it was the case in the early years of the online price comparison industry. In the offline market (e.g. newspapers), "per-click" pricing is still infeasible.

optimality condition:

$$\frac{1}{n_j} + \frac{1-\mu_j}{\sigma}\left\{\left|\log(1-\beta)\right|\left[1+\frac{\beta}{1-\beta}\cdot n_j\right] - \frac{\beta}{1-\beta}\right\}(1-\beta)^{n_j}(r-m) - \left|\log(1-\beta)\right| = 0 \quad (1.32)$$

Equation (1.32) implicitly describes the best response of platform j to the number of sellers listed on platform i.[27] Although the reaction functions appear to be rather complex, in the Appendix I show that the problem has a unique symmetric equilibrium, and also provide numerical evidence that no asymmetric equilibria exist. This observation is summarized in Proposition 1.5.

Proposition 1.5. *If sellers and platforms play symmetric strategies, the model without per-click fees has a unique equilibrium. The equilibrium is characterized implicitly by the per-platform number of advertising firms n^* in the expression:*

$$h(n^*) \equiv \left\{\left|\log(1-\beta)\right| + \left|\log(1-\beta)\right|\frac{\beta}{1-\beta}\cdot n^* - \frac{\beta}{1-\beta}\right\}(1-\beta)^{n^*}\left(\frac{r-m}{2\sigma}\right) + \frac{1}{n^*} - \left|\log(1-\beta)\right| = 0 \quad (1.33)$$

and listing fees:

$$a^* = \beta(1-\beta)^{n^*-1}\frac{r-m}{2} \quad (1.34)$$

provided that the total number of sellers satisfies the minimum condition $N \geq n^*$.

Proof. Setting $n_i = n_j$ in the first order condition (1.32) yields the implicit expression in (1.33). When n^* is small, $h(n^*) > 0$, whereas $h(n^*)$ converges to $-\left|\log(1-\beta)\right| < 0$ as $n^* \to \infty$. Since the function is continuous for $n^* > 0$, at least one symmetric equilibrium exists.

The main difficulty is in showing uniqueness, because for some values of the parameters (β, σ, r, m), $h(n^*)$ is increasing on certain intervals. The remaining technical details of the proof are left for the Appendix. □

Comparative statics

Since the number of advertising firms on a platform directly influences the expected surplus of consumers who visit the platform, it is worth looking at the effect of model parameters on equilibrium outcomes. We can do that by inspecting the implicit equilibrium condition (1.33).

The most important difference between the baseline model with per-click fees, and the extension without, is that platforms do not want to list all sellers when they cannot charge for clicks. Limiting the number of sellers also limits competition among them, enabling the platform to (indirectly) extract more surplus from buyers.

[27] The number of sellers on platform i enters through the consumers' platform choice function: $\mu_j = \mu_j(u_i(n_i), u_j(n_j))$.

When buyers' information is more accurate (σ is lower), the equilibrium number of sellers on a platform increases. Intuitively, since an increase in the number of sellers on a platform is easier for buyers to detect over the random noise, comparison sites can count on attracting more buyers with the sign-up of an additional seller, so they will want to have more sellers.[28] The informational friction in buyers' behavior affects the intensity of platform competition similarly with or without per-click fees.

The probability of a buyer-seller match has an ambiguous effect on the size of a platform's seller base in general. Better matching (higher β) means that fewer buyers are left for the seller who charges the reservation price of consumers, and therefore this seller (and all the others) are willing to pay less for being listed. When the willingness to pay for advertising decreases, so should the equilibrium amount of advertising: platforms should list fewer sellers as a result. On the other hand, a higher match probability also increases the effective stakes (expected transaction surplus) of consumers, causing them to make relatively more informed choices between platforms (just as an increase in $r - m$, or a decrease in σ would). This has the opposite effect on platforms: they are motivated to compete more intensely by signing up more sellers. The net outcome depends on the actual parameters.

1.4.2 Fixed participation costs

When all sellers multi-home, as in my baseline model with per-click charges, the only efficiency loss from having two platforms comes from the duplication of fixed costs—of operating a platform, and of listing prices twice—all of which I have so far taken to be zero. It turns out that the lack of a fixed advertising cost for sellers[29] is not an entirely inocuous assumption concerning the outcome of the game for the following reasons.

Let us assume that, in addition to the listing fees, sellers also incur a real cost $\kappa > 0$ for each advertised offer. In this case, a platform will sign up new sellers until the marginal contribution of the last seller to total expected transaction surplus drops down to κ (as opposed to 0, which it never reaches). If the total number of sellers is large enough, such that the Nth seller to a platform increases surplus by less than κ, then even with per-click charges the platform will exclude some sellers. The question is, will platforms perform this exclusion symmetrically?

Not necessarily, because the marginal contribution of a seller to a platform is larger when the platform has more buyers. If one of the platforms has more buyers to start with, it might (depending on the value of the parameters) also sign up more sellers, which will possibly sustain its higher buyer share as an equilibrium outcome. In addition to the symmetric equilibrium, there may thus exist asymmetric equilibria in which the market is tilted towards one of the platforms. The next proposition formalizes the result on the existence of asymmetric equilibria.

Proposition 1.6. *Platform competition with fixed advertising costs of $\kappa > 0$ on the sellers' part always has a symmetric equilibrium, and also has two asymmetric*

[28]Formally, one can see this relationship from (1.33) by remembering that $h(n)$ is strictly decreasing in the neighborhood of n^* (see the Appendix for a proof). Since the term within the curly brackets is positive, a decrease in σ (or an increase in $r - m$) shifts $h(n)$ upwards. Equilibrium is re-established at a higher value of n^*.

[29]Such as updating prices and inventories in the platforms' database systems, for example.

10.14754/CEU.2015.04

equilibria if the fixed costs are in the range:

$$\sigma\left|\log\left(1-\beta\right)\right| < \kappa < \frac{3}{2}\sigma\left|\log\left(1-\beta\right)\right|$$

For given values of σ and β, the equilibrium is more asymmetric when κ is closer to its lower boundary.

Proof. The derivations are left for the Appendix. □

Following the reasoning above, we can think of a strong asymmetry as market tipping. However, since price advertising markets generally involve several platforms with comparable sizes, the existence of strongly asymmetric equilibria seems to be a theoretical possibility that is not confirmed in practice.

1.5 Discussion

Both buyers and sellers prefer a platform that has a larger number of potential counterparties, holding everything else fixed. Yet, there is no endogenous market concentration in equilibrium. In fact, the two platforms share the market equally. What are the mechanisms that lead to this outcome?

One contributing factor is that I have put a certain structure on how consumers choose which platform to visit. Effectively, this assumption precludes the existence of "bad equilibria" that characterize many two-sided market models. In a "bad equilibrium", both buyers and sellers have the self-fulfilling belief that all members of the other party will show up at the same platform, and hence the other platform gets no business at all.

Instead of allowing buyers to form any beliefs, I introduce a noise term (σ) into their decision making. Save for the degenerate case of $\sigma = 0$, sellers can always be assured that at least a few buyers will show up at each platform. This is enough for offers to be posted everywhere, ruling out equilibria in which one platform is driven out because everyone's beliefs turned out to be against it staying in business.

However, the existence of such degenerate equilibria is purely of technical interest, and hence their exclusion is not a substantial restriction in the model. Market tipping, as understood from an economic perspective, is about the decisive shift in consumer traffic towards one company from among two or more already established market players, brought about by "normal" market forces, rather than an inexplicable coordinated change in beliefs.[30] In this light, an equilibrium in which one platform serves an overwhelming share of the market can also be considered a tipped market, and such asymmetric equilibria are not ruled out by the structure I put on consumer behavior. The central result of the baseline model is that no asymmetric equilibria exist, even when informational frictions are tiny.

The main reason why market tipping equilibria are absent is that all sellers list prices on both platforms. The lack of exclusivity means that no platform can gain

[30] With the introduction of fixed costs to platform operation, one might easily imagine that a small enough market share—even if it satisfies conditions for a short-run equilibrium—will also lead to the dominated platform exiting the market.

10.14754/CEU.2015.04

an advantage over the other on the seller side. If a comparison site finds it profitable, it can always have just as many listed offers as its rival, and hence offer the same utility to consumers. Platforms are therefore only competing on the buyer side, but without the self-reinforcing cycle of network effects: attracting more buyers does not mean that more sellers will also be attracted, because all sellers are already listed. Competition for buyers simplifies to a—slightly modified—symmetric Cournot-game in offering utilities, with the usual market-sharing equilibrium.

This reasoning suggests that we would have to look for signs of market tipping on platform markets on which both sides are either unable or unwilling (because of high participation costs) to multi-home. Consumer auction markets are probably more prone to tipping, for example, because a large proportion of the items for sale are unique, or in limited supply, and the platform conducts the transaction itself. Individual items can easily be advertised, but impossible to be sold, on two marketplaces at the same time.

In the model extension without per-click charges, I also get a unique symmetric equilibrium, but for different reasons than before. When per-click fees are not available, platforms have a single instrument, the listing fee, to regulate both the total surplus generated on the site, and the share of this surplus going to the platform as advertising revenue. These two objectives are opposing and their balancing yields a different type of equilibrium: instead of listing all sellers, platforms will strongly restrict the size of their seller bases.

The restriction of seller-to-seller competition is the result of the strongly diminishing willingness of sellers to pay for advertising as the supply side of the price comparison platform grows. When an additional rival appears on a site, each seller loses β fraction of its formerly expected revenues, *ceteris paribus*. This loss is compensated to some extent by the increased number of visitors to the site, but consumers are imperfectly mobile—and the total consumer base is fixed—so losses will inevitably overwhelm the demand-side gains as the number of advertising sellers grows. The close-to-exponential decrease in willingness to pay for advertising makes total advertising demand elastic, which limits the optimal size of the supply side. Competition on the seller side of a platform neutralizes the potentially self-reinforcing network effect.

One prediction of the platform competition model is therefore that price lists will be relatively long when per-click fees can be charged, and relatively short when they cannot. Empirically, this prediction is difficult to verify, since price comparison sites listing homogeneous products seem to all charge for clicks nowadays (Levin (2011)). Perhaps the main supporting fact of the model's applicability is the ubiquity of per-click pricing, as it also yields more profit for price comparison platforms.

1.6 Related literature

The paper is most closely related to the industrial organization literature addressing the tipping of platform markets towards a monopoly. Ellison et al. (2004) ask whether consumer auction platforms such as eBay and Yahoo Auctions can coexist on a market, or one would inevitably drive out the other over time. They write down a two-sided auction model in which both buyers and sellers prefer having

10.14754/CEU.2015.04

more agents on the other side of the market, but fewer agents on their own side to compete with. It turns out that the two forces balance out (even as the number of participants rises without a bound), leading to a range of equilibria (a "plateau") with two active auction platforms, and hence no necessary market tipping.[31]

My model operates in a different setting (posted prices, as opposed to auctions), but the underlying motivations of the agents are similar. Sellers earn less revenue on a platform when they face more competition, and would therefore prefer another platform with fewer rivals. This effect is fully internalized by the platforms who set prices to maximize their profits, which provides an additional layer compared to the baseline model of Ellison et al. (2004).[32]

There are other differences as well beyond the assumptions on platform behavior. For example, I allow sellers to multi-home (a realistic feature in my setting), which helps the coexistence of platforms, since they only compete on the buyer side of the market. Auction platforms, on the other hand, are not well-suited for multi-homing.

A second line of literature that the chapter contributes to is on search models with an "information clearinghouse".[33] The main feature of these search models is that consumers can access additional price quotes at zero marginal costs, as in Varian (1980). The seminal paper relating clearinghouse-type search models to price comparison platforms is Baye and Morgan (2001). They analyze the price setting decision of a monopoly platform that charges participation fees only, and show that it is optimal for the platform to set low (perhaps zero) fees for buyers to maximize participation, and make most of the profits on higher fees set for sellers. In a subsequent paper, Baye et al. (2011) introduce per-click fees and conclude that a monopoly platform is better-off charging sellers for clicks, rather than for the listing. This observation possibly explains why per-click charges have gradually replaced listing fees in the comparison shopping industry.[34]

The main contribution of my paper to the clearinghouse-type search literature is the analysis of platform competition. I show that platforms compete for buyers by adjusting per-click fees, whereas listing fees are only used to extract surplus from sellers. Per-click fees are lower when buyers make more informed choices about which platform to visit, whereas listing fees move in the opposite direction. Moreover, the equilibrium market structure tends towards equal shares on the single-homing consumer side and endogenously emerging multi-homing by sellers. Hence the market

[31]See also Ellison and Fudenberg (2003) for a more general treatment that is not specific to auction platforms. On the empirical side, Brown and Morgan (2009) have questioned the applicability of the model to the eBay vs. Yahoo Auctions setting. Using field experiments on both auction platforms, they have shown that the main prediction of Ellison et al. (2004) concerning the equality of prices and buyer-to-seller ratios across coexisting platforms does not hold empirically, and hence the market is likely to be in the process of tipping (towards eBay in the U.S.)

[32]The two-sided market model by Hałaburda and Piskorski (2013) also exploits the presence of negative intra-group effects to show that several platforms with different business models might coexist in the same market. Their results rely on the assumption that agents' outside options are different, and agents with worse alternatives are therefore more keen to avoid competing with one another. In exchange, unlike the agents with good outside options, they tolerate a restricted set of choices on a platform. This creates opportunities for platforms to segment the market.

[33]See Baye et al. (2006) for an extensive review.

[34]In the same line of literature, Galeotti and Moraga-González (2009) extend the model of Baye and Morgan (2001) by introducing horizontal product differentiation. Moraga-González and Wildenbeest (2012) analyze the behavior of the monopoly platform with and without the ability to price discriminate between online and offline buyers, and also allow for vertically differentiated products.

10.14754/CEU.2015.04

tipping to a monopoly price comparison site is not a serious issue.

Finally, the topic of the chapter also fits into a large literature on two-sided markets, pioneered by Caillaud and Jullien (2003), Rochet and Tirole (2003), and Armstrong (2006). The form of seller-to-seller competition, as well as the behavior of buyers concerning platform choice, are specific in my model, which yields additional results compared to general two-sided frameworks.

One notable instance concerns a hypthesis by Armstrong (2006) regarding platform pricing when intra-platform (i.e. seller-to-seller) competition is present. Armstrong suggests that platforms will allow or restrict competition depending on whether they can charge consumers for entry. My model confirms this insight with a twist: although buyers are never charged for participation directly, the effect of per-click fees on the distribution of prices is equivalent to making the buyer side pay. When per-click fees can be levied, platforms are motivated to list all available sellers and allow them to compete intensely, whereas with listing fees only, the seller base will be restricted, because there is no other way for the platform to extract surplus from consumers.[35]

1.7 Conclusion

I have built a model of competing information intermediaries (price comparison platforms) between buyers and sellers to investigate whether the intermediation market is prone to tipping towards a monopoly provider. A central feature of my model is that consumers can investigate offers in the order of increasing prices, which leads to a mixed strategy pricing on the sellers' part, and hence to price dispersion on the comparison site, as in the monopoly platform model of Baye and Morgan (2001).

I found that the ability of platforms to charge sellers for consumer traffic ("clicks") fundamentally alters the nature of market equilibrium. When only listing fees are available, platforms will restrict the number of sellers by charging high listing fees to prevent them from competing away the surplus of the buyer-seller transactions. When per-click pricing is feasible, platforms set the listing fee low to attract all sellers and hence maximize total surplus, and use the per-click fee to increase the effective marginal cost of the product, below which no seller sets prices.

I show that the model has a unique symmetric equilibrium both with and without per-click fees, in line with the empirical observation that price comparison markets usually have several players, often with comparably-sized user bases. Market tipping should therefore not be a serious concern for competition policy in this sector.

Two forces act against the cross-market externalities in the price comparison industry: the ability of sellers to be present on more than one platform at a time (multi-homing), and the informational friction with which buyers choose platforms for investigating prices. Multi-homing is the more important feature, since it enables platforms to always copy each other's range of offers, shutting down the seller-to-buyer side externality channel. The informational friction, on the other hand, mainly serves to restrict beliefs about the behavior of the other side of the market, and thus

[35]Belleflamme and Toulemonde (2009) also look at negative intra-group externalities in a general two-sided framework and conclude that they might hinder the entry of new competitors next to an already established free platform.

10.14754/CEU.2015.04

rule out equilibria in which all sellers and buyers coordinate on the same platform regardless of what platforms do.

Online markets often involve activities that scale easily to millions of users, making endogenous market tipping towards a monopoly structure a more prominent competition policy concern than it is in the "offline" economy. This is especially the case when two-sided market externalities are present. In this chapter, I have pointed to opposing market forces that prevent market tipping, and built a model with price-related information intermediaries that sheds some light on how these forces work. Since there are many other types of online intermediation activities, some of which developed towards more concentrated market structures (e.g. internet search or social networks), there are certainly more topics to be explored in this area.

10.14754/CEU.2015.04

Chapter 2

Salience and Switching

2.1 Introduction

People often stick to expensive utility companies, banks, insurers, or other service providers when they would be financially better off switching to an alternative. This inertia can be due to the time or the effort cost of switching, but psychological factors, such as inattention, procrastination, or fear of new situations, can also create or heighten barriers to switching. Measuring the relative importance of these factors is essential for understanding a number of economic issues, from default effects to firm pricing and contracting, as well as for informing competition and consumer protection policy.

If, for example, many people remain in high-priced contracts simply because they do not pay attention to the possibility of switching, then an awareness-raising media campaign can be a cost-effective policy tool for helping consumers and strengthening competition. The same policy, however, will not work if consumers are reluctant to switch due to high effort costs.[1] Similarly, restrictions on contracting between firms and consumers can lead to welfare gains when too many people are inattentive to the "small print", but may be welfare-decreasing otherwise.

In this chapter, I exploit a change in auto liability insurance regulation in Hungary to identify the causal effect of a media campaign that provides no decision-relevant information to consumers, but increases the *salience* of the switching opportunity for a well-defined time period. My main result is that the campaign raises switching rates by 12 percentage points from a baseline of 20 percent. In comparison, the estimated reduced-form relationship between financial incentives and switching decisions is much weaker: an additional saving of $50 per year—or about one-third of the median annual premium—is associated with only 4 percentage points higher switching rates.[2]

The essence of the regulatory change that I exploit is the following. Up to 2009

[1] The influence of switching costs on competition has long been recognized in both academic and policy discussions. For an extensive survey of the switching cost literature, see Farrell and Klemperer (2007). The prevailing view—especially in policy circles—is that a reduction in switching costs tends to intensify competition and increase consumer surplus. Viard (2007) estimates, for example, that the introduction of 800-number portability in the early 1990's has reduced the price of having a toll-free telephone number by 14% in the U.S.

[2] $50 is approximately the median saving in the sample, worth about one and a half days' of average net wages.

10.14754/CEU.2015.04

(in the "old regime"), insurance periods had to coincide with calendar years, and drivers could only switch insurance providers in November, prompting insurance companies and intermediaries to concentrate their marketing activity to the same month as well. In addition, reports on the switching campaign ranked high among the news, raising people's awareness to the issue even further. Yet, the information content of advertisements was severely limited: insurance fee schedules, for example, are much too complex to convey in any type of media message.

Starting from 2010 (in the "new regime"), people who sign a new contract because they buy a car are treated differently. Their once-a-year switching periods are no longer synchronized to November, but remain attached to the anniversaries of their car purchasing dates. At the same time, the gradual shift to the new regime ensures that most switching decisions will still be taken in November and the campaign will live on for several years after the regulatory change. As a result, the switching periods of new-regime drivers will overlap with the campaign if they buy cars close to January 1, but not if they do so in the middle of the year. Since new-regime drivers are otherwise similar to each other, the differences in switching rates—conditional on financial savings and observed individual characteristics—must arise from being close to, or far from, the campaign period.

For the estimation, I collected data from an independent intermediary in the Hungarian auto liability insurance market covering the years 2009-2012. The dataset includes contract-level information about insuree demographics, vehicle characteristics, payment options, the identity of the insurer, the insurance fee for the first year, and the start and end dates of the contract, from which the act of switching can be deduced. I calculate all available price offers in the market for each person using the public price schedules of insurance companies, and define the financial savings from switching as the difference between the continuation price of the existing contract and the cheapest alternative offer. In most specifications, campaign treatment is a dummy variable that indicates whether a person's time window for contract switching overlaps with the media campaign in November for at least one day, although I also explore how switching rates evolve as time passes between campaigns.

I measure the effect of the campaign on switching decisions in a simple discrete choice model.[3] The dependent variable is whether a person switched contracts at the end of his first year. The two main explanatory variables are the campaign dummy and the amount of financial savings gained by switching. I also include various interaction terms in the regression and control for all observable characteristics.

The results show that salience has a large effect on switching. The baseline switching rate is 20 percent, which increases by 12 percentage points during the campaign. Financial incentives seem to matter less: the difference in switching rates between people whose savings are at the 90th and at the 10th percentile ($12 and $110, respectively) is only 8 percentage points. Moreover, the campaign has a proportionally larger effect on low savers, suggesting that ignorance of the switching opportunity is not purely by chance: it is those who stand to gain less on average who need to be reminded by the campaign to shop around. Despite this hint of rationality, the large overall campaign estimate still shows that inattentiveness to

[3] I employ a logit framework, but the results are robust to other functional forms (probit, linear probability) as well.

10.14754/CEU.2015.04

switching is suboptimal for many people.

From a policy perspective, my results indicate that consumers could derive considerable benefits from effective information-spreading and market education campaigns, as well as a market design that makes infrequent, but economically significant choice situations more salient.[4] In particular, there are a number of important consumer markets with low switching rates and weak competition (e.g. gas, electricity, banking) in which consumers could benefit from an endogenously arising campaign effect if switching opportunities were restricted to specific times of the year.

This is the first paper to measure the causal effect of inattention on consumer inertia using a micro-level observational dataset. More generally, the effect of inattention on choices has been studied empirically in a handful of recent papers. Chetty et al. (2009) have looked at the consequences of posting after-tax prices in addition to the usual pre-tax prices in a supermarket experiment. Although they demonstrate with surveys that people are well aware of the sales tax they have to pay at the counter, making it salient at the time they pick the items off the shelf still has substantial negative demand effects. Their experimental conclusions are backed by the analysis of excise tax and consumption changes of alcoholic beverages in observational data.

Finkelstein (2009) has studied the introduction of electronically paid highway tolls in various U.S. states over time, and has found that fees were set 20-40 percent higher whenever the payment was deducted automatically, as opposed to being handed over in cash at the toll booth. She attributes the demand differences to the salience of prices at the time of purchase.

Hossain and Morgan (2006) have carried out an online auction experiment and found that buyers do not take into account the (less salient) shipping fees in a fully rational manner when bidding for a product. In particular, higher—but non-excessive—shipping fees increased the total auction revenue for some products. Einav et al. (2014) have qualitatively confirmed this result in a large observational dataset.

In financial markets, DellaVigna and Pollet (2009) and Hirshleifer et al. (2009) have studied the differential effects of earnings announcements on various days of the week. They have found that announcements on Fridays, or on days that are "congested" with news, take longer to be incorporated into the price of assets. This result also suggests that investors' limited attention has measurable effects on asset price dynamics.

The rest of the chapter is organized as follows. Section 2 provides a background on the vehicle liability insurance market in Hungary. In section 3, I describe the data and the identification of the campaign effect, and substantiate the comparability of people in the distinct treatment groups. In section 4, I estimate the discrete-choice specifications for contract switching, followed by a discussion of the chapter's

[4]The direction of the regulatory change in the insurance market I study is the opposite: it makes consumers less, rather than more, aware of the choices they face. It is, therefore, a step backwards from a consumer policy point of view. From an overall welfare perspective, which also accounts for firms' profits, the picture is less clear. The co-existence of attentive "switchers" and inattentive "non-switchers" in the market can lead to mixed-strategy equilibria in pricing (Varian (1980), Baye and Morgan (2001)), or a dynamic pattern in which high- and low-pricing firms change roles in each period (Farrell and Shapiro (1988)). In such a situation, increasing the share of switchers is welfare-decreasing, even if it benefits consumers.

10.14754/CEU.2015.04

results and concluding remarks. Details about the preparation of the dataset, the calculation of alternative insurance premia, and robustness checks are left for the appendix.

2.2 The vehicle liability insurance market in Hungary

2.2.1 General rules

Auto liability insurance is a mandatory product, vehicles are not allowed to participate in road traffic without it. The insurance covers property damage and bodily injury to third parties only, in case the driver of an insured vehicle is found to be at fault in an accident.

Coverage is provided by over a dozen insurance companies and one mutual insurance association. The terms of the service are regulated by law, but companies are free to set their own insurance premia. Prices are set once for each calendar year, and cannot be changed during the year. The time for announcing fees for the upcoming year is at the end of October.[5]

The usual length of an insurance contract is one year. Upon expiry, motorists are free to switch insurance companies, but only if they send a cancellation note to their existing insurer 30 days before the next insurance period starts. If they failed to provide notice in time, their existing contract will be automatically renewed for another year at the continuation price set by the insurer.

Continuation prices are announced to clients about 2 months before the existing contract runs out. These prices might be different from the price offered by the same company on a new contract with identical characteristics. People can also choose a new contract with their existing insurer instead of continuing the old one (*re-contracting*), but the same 30-day cancellation note rule applies in this case as well.

Contract switching prior to 2010

Until January 1, 2010 (in the *old regime*), contracts were required to coincide with the calendar year from the second year onwards. If, for example, someone bought a car in July 2008, his first insurance period only ran until December 31 of that year, and the second and subsequent periods covered the years 2009, 2010, etc. in full.

By the rules of notification, drivers in the old regulatory regime had to cancel their existing contracts by December 1st if they wanted a different one for the next calendar year. Since prices were announced at the end of October, people had the month of November to consider changing insurance contracts. Figure 2.1 shows the timing of events before 2010.

Insurance companies and intermediaries used this synchronized switching opportunity to advertise heavily to consumers. According to industry sources, market players spent at least 90 percent of their yearly marketing budgets in November. The switching campaign was also regularly covered in various media outlets.[6]

[5]The once-a-year price setting restriction has been removed from the regulation starting from January 2013. The

Figure 2.1: Time periods for switching were synchronized for all insurees before 2010

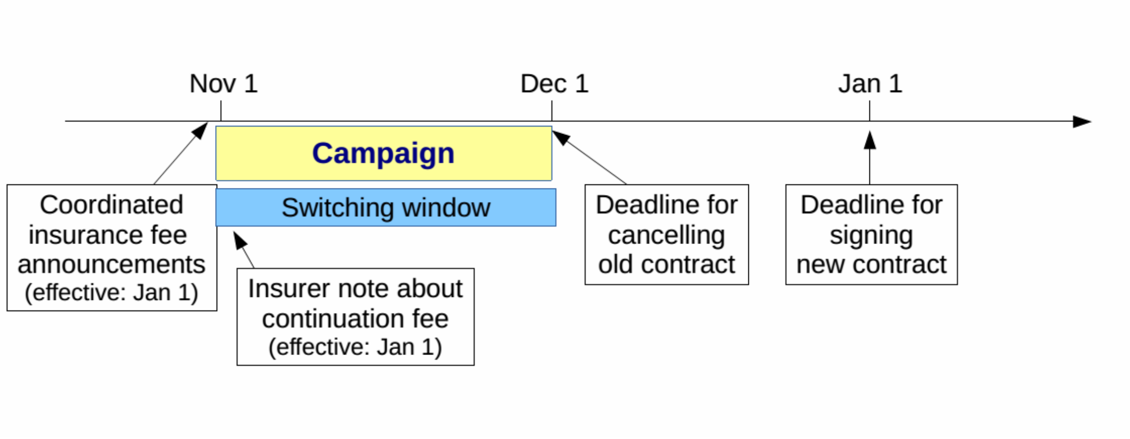

Contract switching since 2010

Following January 1st, 2010, all new insurance periods—including the first one—have become one year long. As a result, if someone bought a car in the middle of May 2010, the first time he could switch contracts was not in November 2010, but between mid-March and mid-April in 2011 (see Figure 2.2). The other elements of the system (once-a-year price announcements, 30-60 days notification rules, etc.) remained unchanged for the time period that I study.

The move to the new system is gradual, in that it only affects people after a vehicle acquisition. If a person already had a contract on January 1, 2010, and has not changed vehicles since then, then his opportunity to switch is still in November every year.

In the initial years of the new regime, the switching campaign in November was as intense as before. Prices were still announced at the end of October, and most people (those who did not change cars since 1/1/2010) were still end-of-year switchers.[7]

2.2.2 Structure and calculation of the insurance premia

Risk categorization

People's driving histories are recorded according to a regulated classification method, called the *bonus-malus system*. The main idea is to reward those who have a longer accident-free past with insurance fee reductions, and to increase the comparability of offers across insurers.

focus of this chapter is on years 2010 and 2011.

[6]The campaign had directly measurable effects on the salience of the switching opportunity in consumers' perception. A survey, commissioned by the Hungarian Competiton Authority in 2009, found that almost 30 percent of the respondents considered the switching decision because of hearing about it during the campaign. Another 30 percent claimed that they would have shopped around with or without the campaign, whereas the rest were totally ignorant of the switching opportunity (Scale Research (2010)).

[7]The share of Calendar clients in the insured population decreases by about 10 percentage points in a year. Eventually, as most people change vehicles, contracting dates will be dispersed evenly throughout the year.

10.14754/CEU.2015.04

Figure 2.2: Time periods for switching are individual-specific since 2010

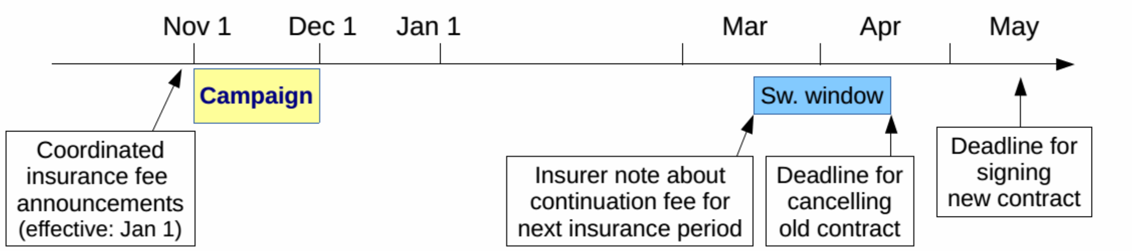

For passenger vehicles, the system contains 15 categories: M4–M1, A0, and B1–B10, in increasing order. Everyone starts in A0. Driving for one year without causing an accident increases one's rating by 1 (from A0 to B1, say), and causing an accident decreases it by 2 (from A0 to M2, for example).

Ratings above B10, or below M4, are not allowed. If a person has been uninsured for at least 2 years, or already has an insurance on another vehicle, his rating can only be A0. Ratings characterize the drivers, not the vehicles, and as such they are transferable across cars and across insurance companies.

Although price structures are complex, risk ratings generally enter the pricing formulae as multiplicative terms. Drivers in B10 can expect to pay about half as much as those in A0, whereas the penalty factors in M4 raise insurance premia by 100-300% relative to A0. Companies are free to set their own multipliers, but cannot deviate from the system itself.

Other pricing factors

Insurance companies use a host of criteria to discriminate between drivers in their pricing. The most common are the age and home address of the driver, the power and usage of the vehicle, and the frequency and method of payment. You pay more if you are younger (especially under 35), live in a larger town, have a more powerful car, use your vehicle for non-personal purposes (e.g. as a taxi), pay in monthly or quarterly (rather than yearly) installments, and pay by check (rather than wire transfer).

These pricing factors are complemented by several other discounts or penalties, the use of which varies widely across companies. For example, you might have a different insurance premium if you are a returning client, accept electronic communication means, have children (but not of driving age), have been driving for a longer time, buy other types of insurance products from the same company, have a less common car brand, and so on.

Insurance companies also employ a number of different methods to calculate the final premium. Some insurers have tables with basic fees, which are multiplied by the applicable discount factors and the risk rating. Others have a score-system

10.14754/CEU.2015.04

in which one first calculates the applicable discount or penalty scores, adds them together, then looks up the corresponding basic fees from scoring tables. Still others combine additive and multiplicative discounts, use sets of discounts of which only one can be chosen, or put caps on the overall discount relative to the basic fee.

Availability of price information

Insurance companies are required to announce their pricing rules and tables publicly at least 60 days before they take effect. Before 2013, price setting was restricted to calendar years by law, and hence the announcements were made on the last day of October, in the form of coordinated advertisements in two national newspapers.[8] In addition, own price tables and methodology are available on each company's website.

The price publicity itself does not count for much, however, since it would take an exceedingly long time for anyone to calculate their insurance premia at all the companies on the market, due to the complexity and variability of the price setting methods. Even then, there would be considerable uncertainty left whether one has applied all the pricing rules correctly.[9]

There are two alternative routes for shopping around. The first one is to visit the local offices (or websites) of insurers, provide the necessary information for price setting, and let the company staff (or server) carry out the calculations. Since there are over a dozen companies on the market, this is still a very time-consuming approach.

The second route involves the use of insurance brokers, many of whom have both online and offline presence. They operate free price comparison engines, so that one only needs to provide information once to see all quotes on one page. If authorized, the brokers will also take care of all administrative tasks, including the cancellation of existing contracts and the signing of the new ones. For this service, they earn a percentage cut from the insurance premia, paid by the insurers. The use of insurance brokers, therefore, provides car owners with a relatively cheap way of shopping around. All drivers in my dataset used an insurance broker when they originally signed their contracts.

2.3 Data and identification

I collected data from a mid-sized insurance broker company in Hungary for the empirical analysis.[10] To facilitate contracting between drivers and insurers, the company keeps a record of all personal and vehicle characteristics relevant to a liability insurance contract, as well as general contract features. Since brokers are paid yearly commissions for a contract by the insurance companies, these records also form the basis of financial settlement.

[8]The complexity of the fee structures is well-illustrated by the fact that, in some years, the ads of 15 companies filled a 140-page long attachment to the dailies.

[9]For some insurers, the public announcement of price setting rules is reminiscent of the small print of legal contracts, as if the company did not actually expect consumers to use the announcement itself to calculate prices.

[10]Drivers are free to choose among a number of insurance brokers, as well as contact the insurers directly. Therefore, my dataset has no claims to representativeness across the population. In particular, long-time owners who rarely switch insurance contracts are likely to be under-represented in the sample.

Figure 2.3: Identification of the campaign effect relies on the exogenous timing of individual switching periods

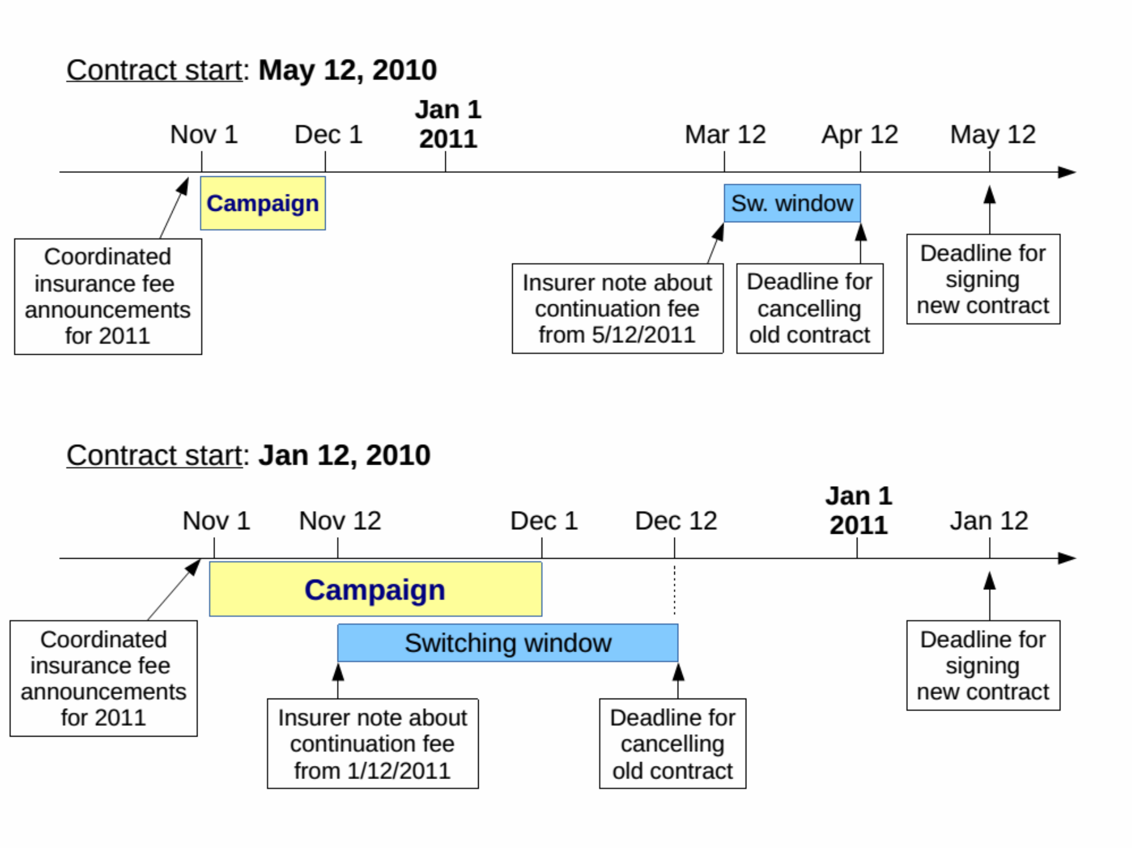

The client base of my broker company consists of two main sources: online and dealerships. The online interface is the company's home page, on which drivers can compare and choose insurance offers, initiate contracting, and discontinue existing contracts. Clients are also acquired from hundreds of car dealerships across the country where the broker company has representatives to take care of liability insurance right after a vehicle purchase. To compare offers at the dealership, the representatives also use the company's online interface. In subsequent years, the drivers can use the broker's online price comparison tool themselves to switch contracts.

In addition, the company has a number of affiliates scattered across the country and one customer center in a large city, but both of these are relatively minor sources of new clients.

2.3.1 Identification

I identify the causal effect of the campaign on switching rates by comparing the switching decisions of people who are close in time to the campaign with the decisions of those who are farther away. Figure 2.3 provides a detailed example.

The top panel of Figure 2.3 shows a contract with a starting date of May 12,

10.14754/CEU.2015.04

2010. The first switching period for this contract is between March 12 and April 12, 2011, which is more than three months after the campaign of November 2010.

In the bottom panel of the same figure, a contract starts on January 12, 2010. Its first switching period is between November 12 and December 12, 2010, which largely coincides with the campaign at the end of 2010.

My conjecture is that the switching rates at the first anniversaries of these two contracts will be different, and the difference will be attributable to the switching campaign itself.

In identifying the causal effect of the campaign on switching rates, my main assumption is that the unobserved determinants of switching are mean-independent of a contract's starting date between January 2nd and December 31st in 2010, once we control for contract prices and the observable characteristics of drivers.

The most common reason for starting a liability insurance contract in 2010 (except for 1/1/2010) is the acquisition of a vehicle. I assume that people do not purposefully time this decision to January (instead of, say, June) in order to benefit from the "social reminder" mechanism of the switching campaign.

There are three strong arguments why this assumption is a valid one. First, the financial stakes involved in a car and an insurance buying decision are about two orders of magnitude apart. A potential 20 percent loss on a liability insurance that should have been cancelled in time is still small change when compared to the price of a vehicle.

Second, if people were conscious about liability insurance pricing, they should time their car purchasing decision to January 1st, since many insurers explicitly target this group with extra discounts (worth about 10-15 percent of the baseline price).[11] In the data, no one actually bought a car on January 1, 2010.

Finally, we know from psychological research that people are often naive about their cognitive limitations, especially when future scenarios are concerned. It is very unlikely that they would take costly precautions (such as delaying car purchases) to avoid forgetting a "trivial" decision like changing an insurance contract.

Laying aside the conscious timing of vehicle acquisition, there might also be other ways in which the identifying assumption might fail. It is possible, for example, that people buying cars at the beginning of the year are somehow different from people who buy later, and this difference matters for switching.[12]

These objections to the identification argument can best be countered by examining Figure 2.4. In the figure, the horizontal axis is a timeline of contract starting times between January 1, 2009 and December 31, 2011. The vertical axis measures switching rates. Each dot in the graph represents the proportion of contracts that were cancelled at the first switching opportunity within all the contracts started on the given week. I only included insurance contracts that were signed because of a vehicle acquisition.

[11] The reason for the January 1st discounts is that most existing contracts had been signed in the old regime, and therefore turn on the first day of the year. In addition to the sheer number of contracts, this group contains many people who have switched contracts before, which indicates that they are more sensitive to prices.

[12] For example, there might be unobserved car discounts in January, which attract a larger proportion of financially savvy consumers. In theory, differential macroeconomic effects within the year (e.g. changing economic outlook) might also play a role in the selection of car buyers. In practice, however, I find no evidence for changes in macroeconomic conditions.

Figure 2.4: Switching rates at the end of the first insurance period are consistently higher during campaign periods in both the old and the new regimes

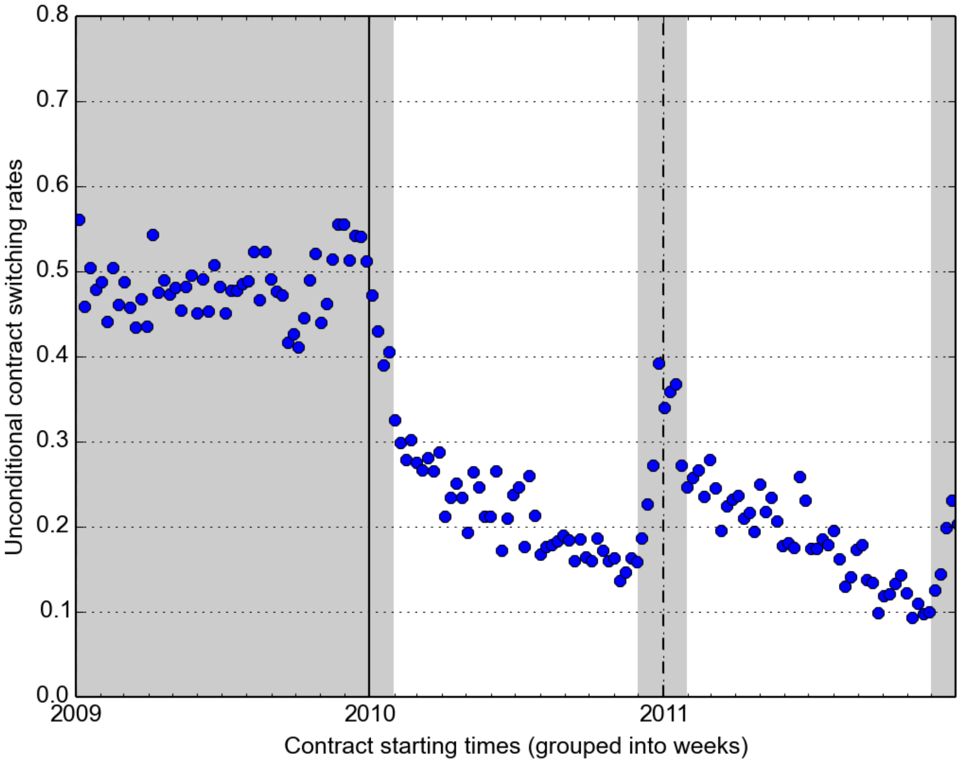

The shaded area in the figure covers the contracts whose switching period overlapped with an advertising campaign in November. This includes all contracts in the old regime (2009), since they were synchronized to January 1, as well as December and January contracts in the new regime (2010-2011). The date of the regime change (1/1/2010) is marked by a solid vertical line in the figure.

The identification strategy is validated by the observation that switching rates of January clients are not different from those of non-January clients in the old regime, when everyone was equally affected by the campaign. If the identification assumption failed (perhaps because January clients were more financially savvy), we would expect to see at least some difference in switching among the 2009 contracts.

The unconditional numbers in Figure 2.4 also show that the campaign did make a difference in the new regime. Switching rates for January 2010 contracts are high, whereas the rate declines from February to November, only to rise again towards the end of the year. The pattern repeats itself in 2011, proving that we are not witnessing a simple downward trend in consumers' willingness to shop around.

10.14754/CEU.2015.04

2.3.2 Descriptive statistics

My baseline dataset contains around 14,000 contracts that were started in 2010 and were still in effect at the first anniversary. This is a restricted sample,[13] because I only include contracts for which I can accurately calculate the insurance fee that I observe in 2010 (within $\pm 5\%$). The reason for the restriction is to limit the measurement error introduced by the calculation of unobserved alternative insurance premia (see more about sample construction in the data appendix).

45 percent of the baseline dataset (the *Calendar* group) is made up of old-regime drivers who have had liability insurance on their current car for a while, and decided to switch contracts on January 1, 2010. These people are likely to be different along several unobserved dimensions from the drivers who are not affected by the campaign. Therefore, the *Calendar* group will not play a role in identifying the campaign effect, and I will only include its descriptive statistics whenever doing so puts the comparison of the other groups into better perspective.

The non-*Calendar* contracts are distributed fairly evenly throughout 2010, but with a noticeable upward trend towards the end of the year. The lowest monthly number is 326 (January 2nd-31st), and the highest is 831 (November 1-30th).[14]

Besides the *Calendar* group, I also drop the contracts started in December 2010 from the baseline estimates. Their switching window is so far away from the relevant campaign that it partly overlaps with the next campaign. It is, therefore, hard to tell unequivocally whether they should belong to the treated or the control group. I leave them out of the main specifications, but include them when investigating for separate monthly effects and for robustness.

The remaining sample contains 6,766 contracts, all of which were started between January 2nd and November 30th in 2010 following the acquisition of a vehicle. To determine whether a driver is potentially affected by the switching campaign, I use the following criteria. If the 30-60 days switching window overlaps with the campaign period for at least one day, then the contract is "treated", otherwise it belongs to the control group. The affected contracts are therefore (with a slight rounding) the ones starting in January 2010. For the rest of the chapter, I will use the terms *Campaign* and *No campaign* to denote the treatment and the control groups.[15]

Table 2.1 shows the means and standard errors of the *Campaign* and the *No campaign* group for various observable personal, vehicle, and contract characteristics, as well as a test for differences in the means of the two groups.[16]

The main lesson from Table 2.1 is that the observable personal characteristics and

[13] The original data contains about 2.5 times as many observations as the baseline dataset.

[14] A comparison with aggregate market data on car purchases reveals that this trend is an artifact of the broker company's increasingly successful client acquisition strategy, rather than an increase in overall monthly car purchasing rates.

[15] In reality, treatment status is not as binary as this rule suggests. Even within the *Campaign* group, the switching window of some drivers overlaps with the campaign more than others. Also, the switching deadline does not bring the campaign itself to a full stop (billboards come down gradually, switching outcomes are reported a few days later, etc.) As a result, it's more accurate to say that the exposure intensity to the campaign was high for January drivers and decreased afterwards with a sharp drop in early February. In the extensions, I will study other specifications with more nuanced treatment effects.

[16] Significance of mean differences and estimated coefficients are denoted for the 10%, 5%, and 1% levels in the usual manner in all tables of the chapter.

Table 2.1: Descriptive statistics on insurees, vehicles, and contract characteristics

	Campaign	No campaign	Mean diff.
Age	40.15	40.27	-0.12
	(0.68)	(0.16)	(0.69)
Female	0.298	0.327	-0.030
	(0.025)	(0.006)	(0.026)
Capital resident	0.132	0.136	-0.004
	(0.019)	(0.004)	(0.019)
Years since license obtained	18.29	17.77	0.52
	(0.63)	(0.15)	(0.65)
Mean bonus grade	1.552	1.551	0.001
	(0.181)	(0.041)	(0.185)
Power (kW)	70.16	67.06	3.09**
	(1.44)	(0.31)	(1.48)
Vehicle age (years)	9.095	9.459	-0.364
	(0.279)	(0.066)	(0.286)
Share of premium brands	0.126	0.067	0.058***
	(0.018)	(0.003)	(0.019)
Share of Opels	0.153	0.157	-0.003
	(0.020)	(0.005)	(0.020)
Share of diesel cars	0.319	0.273	0.046*
	(0.026)	(0.006)	(0.026)
Car dealership contracts	0.914	0.926	-0.012
	(0.016)	(0.003)	(0.016)
Payment by check	0.957	0.936	0.021*
	(0.011)	(0.003)	(0.012)
Quarterly payment	0.939	0.940	-0.002
	(0.013)	(0.003)	(0.014)
Insurance premium ($)	193.2	185.5	7.7
	(5.3)	(1.2)	(5.4)
Observations	326	6,440	

10.14754/CEU.2015.04

risk classifications of *Campaign* and *No campaign* drivers are essentially identical. None of the variable means are significantly different from one another at even the 10% level.

Overall, drivers in the *Campaign* group have somewhat larger, more powerful, and more expensive cars than the *No campaign* people.[17] This difference turns out to be a peculiar quirk in sample construction and has no effect on the results of the paper (see the robustness checks and the data appendix for more details).

Most of the contracts were signed at the car dealership following a vehicle purchase. The fees are usually paid in quarterly installments using postal checks (that is, in cash at a post office). There is no difference between *Campaign* and *No campaign* drivers in this respect.

The unconditional point estimate for the mean insurance fee difference between the two groups is $7.4, although it is not statistically significant.[18]

2.3.3 Unobserved heterogeneity across groups

Although I have information on personal, vehicle, and other contracting characteristics, a number of additional pricing factors, such as recent accident history, the buying of additional insurance products, or the length of previous contracts, are not available in my dataset. These unobserved factors might vary systematically across groups, potentially influencing the switching decision.

It might be the case, for example, that *No campaign* clients have more than one insurance product with the same company, which decreases their willingness to switch providers for liability insurance. This effect is unrelated to the salience mechanism of the campaign period, but I would not be able to distinguish the two from the data.

However, several insurers provide discounts on liability insurance if a driver has other insurance products, and I do observe the liability insurance premia. Therefore, price-relevant unobserved heterogeneity across treatment and control groups can be detected by comparing insurance fees conditional on observable characteristics. The importance of this exercise is underlined by the apparent (although not statistically significant) difference in mean insurance prices reported in Table 2.1.

Specifically, I run an OLS regression of observed insurance fees in 2010 on a dummy variable indicating the *Campaign* treatment and include various sets of control variables. The results are shown in Table 2.2.

The first column of the table contains the unconditional price differences, the point estimates for which can also be derived from the "Insurance premium" row of Table 2.1. Controling for an increasing number of observed characteristics initially increases the accuracy of the estimated $7 difference between the *Campaign* and *No campaign* groups, but does not really change the point estimate itself. However, when the identity of the insurer is taken into account (column 5), the measured price difference disappears.

[17] The following brands in the sample are categorized as premium: Audi, BMW, Jaguar, Lexus, Mercedes-Benz, Mini, Porsche. Suzuki is a relatively inexpensive car brand, while Opels are the most common vehicles. Diesel-fuelled cars are generally more expensive than comparably-equiped gasoline-fuelled ones.

[18] Dollar amounts are calculated from Hungarian Forints using the 2010 average exchange rate of 208.15 HUF/USD.

Table 2.2: OLS estimates of insurance premium differences

	(1)	(2)	(3)	(4)	(5)
Campaign	0.077	0.115***	0.074**	0.075**	0.004
	(0.054)	(0.038)	(0.031)	(0.031)	(0.025)
Risk controls		Yes	Yes	Yes	Yes
Personal controls		Yes	Yes	Yes	Yes
Vehicle controls			Yes	Yes	Yes
Payment controls			Yes	Yes	Yes
Contract channel controls				Yes	Yes
Insurer controls					Yes
Observations	6,766	6,766	6,766	6,766	6,766

Notes: The reference group is *No campaign* and all regressions include a constant. Robust standard errors are reported in parentheses. Risk controls: dummy variables for each risk category. Personal controls: age (cubic polynomial), gender, residence in capital. Vehicle controls: power (cubic polynomial), dummies for premium brands and Suzukis, fuel type, vehicle age. Payment controls: dummies for payment frequency and mode. Contract channel controls: online, dealership, or other. Insurer controls: dummies for each company.

Overall, the conclusion from Tables 2.1 and 2.2 is that *Campaign* and *No campaign* drivers are essentially identical in both their observed, and their unobserved but price-relevant, characteristics. This result strengthens the basic identifying assumption of the campaign effect.

2.3.4 Switching rates

My dataset contains the date at which a contract ended, as it was recorded by the broker company. Comparing the start and end dates, liability contracts starting in 2010 can be classified into the following 5 duration categories: (1) less than 1 year; (2) exactly 1 year; (3) more than 1, but less than 2 years; (4) exactly 2 years; and (5) more than 2 years.

A contract can end on a date that is not an anniversary if the vehicle's owner or operator changes, if the vehicle is taken out of traffic, if the insurance company cancels the contract because of non-payment, or if the contract is cancelled by mutual agreement for any reason. Out of these, ownership change is the typical scenario and the rest are relatively rare.

In theory, all of the above reasons might coincide with an anniversary as well, but their low probability on any given day makes it very likely that a contract length of exactly 1 or 2 years implies conscious switching on the driver's part. Since I do not observe the actual reason for a contract's ending, this is what I will assume.[19]

Table 2.3 shows the unconditional switching rates of first and second-year dropouts, as well as first and second-anniversary contract switches. In addition to the *Cam-*

[19]See the data appendix for more details about pinpointing switchers.

10.14754/CEU.2015.04

Table 2.3: Unconditional contract-ending hazard rates

	Campaign	No campaign	Mean diff.	Calendar
Dropout rate in Year 1	0.235	0.197	0.038*	0.181
	(0.021)	(0.004)	(0.021)	(0.004)
Observations	426	8,018		7,908
Switching rate at Ann. 1	0.374	0.203	0.172***	0.469
	(0.027)	(0.005)	(0.027)	(0.006)
Observations	326	6,440		6,476
Dropout rate in Year 2	0.157	0.183	-0.026	0.147
	(0.025)	(0.005)	(0.026)	(0.006)
Observations	204	5,135		3,438
Switching rate at Ann. 2	0.302	0.212	0.090**	0.298
	(0.035)	(0.006)	(0.036)	(0.008)
Observations	172	4,195		2,932

Notes: Each cell contains the estimated probability of a contract's ending within the given time interval, conditional on the contract's "survival" up to that time (number of such contracts also reported). Standard errors are in parentheses.

paign and *No campaign* groups, I also show the corresponding figures for the *Calendar* group for comparison.

Dropout rates are largely similar across the non-*Calendar* groups, and slightly higher than in the *Calendar* group. On the other hand, there are large differences in the switching rates. *Campaign* drivers are almost twice as likely (17 percentage points difference) to switch contracts at the end of the first year than *No campaign* drivers.

At the end of the second year, the switching probability of *No campaign* drivers is about the same as a year earlier, but those of the *Campaign* and *Calendar* drivers have declined considerably.[20] Still, the remaining *Campaign* insurees are more likely to change insurers than the remaining *No campaign* drivers. The mean switching rate differences are highly significant in both years.

I could estimate a campaign effect in the second year as well as the first. The second year's estimate, however, would almost certainly be biased. Since people in the *Campaign* group are more likely to switch at the end of the first year than the *No campaign* drivers, the samples will be differently selected by the end of the second year. In particular, the first-year non-switchers in the *Campaign* group are probably more averse to switching than the first-year non-switchers of the *No campaign* group, which would bias the estimated second-year campaign effect downwards. For this reason, I do not report the estimated second-year effects.

[20]The simplest explanation for the decline is that the people who did not switch at the end of the first year are probably less sensitive to the campaign's "nudging" (they might be more insulated from advertising, for example).

10.14754/CEU.2015.04

Table 2.4: Descriptive statistics on switching benefits ($) in 2011

	Campaign	No campaign	Mean diff.
Mean	62.48	56.37	6.12
Std.err.	3.09	0.61	3.15
Std.dev.	55.76	49.32	
25%	28.93	26.43	
50%	46.29	44.10	
75%	83.51	72.18	
Observations	326	6,440	

2.3.5 Switching benefits

I define the monetary benefit of switching as the difference between the cheapest alternative offer and the default continuation price at the first switching opportunity. This variable captures most accurately what drivers gain by changing contracts.

As I have no data on insurance premia in 2011, I use the pricing schedules of insurers to estimate the prices that individual drivers would have paid, had they chosen any of the alternative offers. The price calculation methodology and the potential measurement issues are discussed in detail in the appendix.

Table 2.4 shows descriptive statistics about the estimated switching benefits by driver group. For example, drivers with a contract anniversary in May 2011 (*No campaign* group) could decrease their insurance premia by $56.6 on average, if they were to cancel their current contract and choose the cheapest available insurance company instead. The benefit distribution is skewed to the left: the savings of the median driver only amount to $44.6 in the same category.

The *Campaign* group has somewhat higher expected benefits from changing insurers than the *No campaign* group. The difference is about $6 and significant at the 5% level. This is the flip side of the approximately $7 difference in the initial insurance fees (see Table 2.1), and is explained by the same peculiarities of sample construction I described earlier.

2.4 Results

2.4.1 Baseline estimates

I estimate the probability of switching contracts in 2011, conditional on not dropping out before the first switching opportunity, using a standard binary logit model. The main variables of interest are the campaign dummy and the calculated monetary benefits of switching. As in Table 2.2, I use several sets of control variables. The results of the estimation are shown in Table 2.5.

The top panel of the table shows the estimated marginal effects for the campaign treatment, the calculated switching benefits (measured in $100), and their interaction. The bottom panel indicates which sets of control variables were included in the various specifications.

Table 2.5: Logit marginal effects on the probability of switching at the first anniversary of the insurance contract

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Campaign	0.120***	0.116***	0.116***	0.117***	0.120***	0.116***	0.119***
	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)
Switching benefit ($100)	0.091***	0.096***	0.097***	0.091***	0.092***	0.091***	0.075***
	(0.009)	(0.009)	(0.010)	(0.011)	(0.011)	(0.011)	(0.013)
Campaign $\times$ Sw. benefit	0.024	0.027	0.028	0.026	0.024	0.027	0.013
	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)	(0.037)
Risk controls		Yes	Yes	Yes	Yes	Yes	Yes
Personal controls			Yes	Yes	Yes	Yes	Yes
Vehicle controls				Yes	Yes	Yes	Yes
Payment controls					Yes	Yes	Yes
Contract channel controls						Yes	Yes
Insurer controls							Yes
Observations	6,766	6,766	6,766	6,766	6,766	6,766	6,766

Notes: All contracts start between January 2 and November 30, 2010.

10.14754/CEU.2015.04

10.14754/CEU.2015.04

The first thing to note is that both the campaign dummy and the switching benefit variable are highly significant in all columns, whereas the interaction term is never different from zero. Moreover, the point estimates are almost identical, regardless of the employed control variables. This suggests that the estimates are robust and not biased by omitted variables.

The estimated marginal effect of the campaign on switching rates is large. Drivers whose switching window coincides with the campaign are 12 percentage points more likely to change insurance contracts than the rest of the sample. The estimates are robust to functional form: linear probability and probit models yield the same results.[21]

The estimated relationship between monetary incentives and switching decisions is much weaker: $100 of additional yearly savings are associated with only 8 percentage points higher switching rates.[22]

Considering that switching benefits are below $150 for almost everyone in the sample,[23] the estimated marginal effect suggests that pecuniary savings have little influence on switching behavior.

There is, however, a different interpretation of the results. According to survey evidence cited earlier, many people—even during the campaign period—are completely ignorant about the opportunity to switch contracts. Their presence in the sample will necessarily bias the price sensitivity estimates towards zero.

Besides a large campaign effect and a weak relationship between savings and switching, I estimate a slightly positive, but statistically insignificant interaction coefficient between the campaign dummy and the benefits of switching. This result implies that the campaign raises the switching rates of insurees having high and low savings with equal *percentage points*. Since people with lower savings switch less often, the campaign actually increases the share of low-saving insurees among all the switchers. Ignorance of the switching opportunity is therefore not purely by chance: it is those who stand to gain less on average who need to be reminded by the campaign to shop around.

2.4.2 Treatment effects by month

The *Campaign / No campaign* treatment cutoff is not necessarily binary. In this section, I explore the implications of a finer treatment effect structure.

One could think of the January-February, or January-March, difference as the "primary effect" of the campaign, in the sense that drivers with contract anniversaries in January are fully subjected to the campaign during their switching window, whereas drivers in March are not subjected to it at all.[24]

[21] In a linear probability model, the specification in the first column of Table 2.5 yields a campaign coefficient of 0.172, exactly the same as the corresponding unconditional mean difference in Table 2.3. When control variables are also included, the linear probability estimates are essentially the same as those reported in Table 2.5.

[22] I get the same results by separating the switching benefit into its two components: the continuation price and the cheapest alternative price. The point estimates for the two prices are statistically equal in absolute terms (close to 0.08), but have opposing signs.

[23] The 95th percentile of the benefit distribution is at $148.

[24] Arguably, early February drivers could belong to the *Campaign* group for reasons mentioned before. It is therefore less ambiguous to compare the first and the third months of the year, rather than the first and the second.

However, drivers with contract dates in any other month of the year also live through the switching campaign. Only, they do not immediately act on the "nudging" to shop around for insurance. It is nevertheless conceivable that the campaign messages are remembered later during the year to some extent, which we can label the "secondary effect" of the campaign. It is also plausible that the secondary effect would die down over time as people think about it less and less. Figure 2.4 confirms this intuition graphically.

For a more accurate quantification, I re-ran the logit regressions of the binary treatment case, but substituted the *Campaign* dummy with monthly fixed effects.[25] I switched the reference group to January for presentation purposes. The results are shown in Table 2.6.

All coefficients are significant and stable across the different specifications.[26] The primary effect of the campaign is an increase in switching rates by about 5-7 percentage points (January-February or January-March comparison). The secondary effect accumulates gradually over the next 8 months (following March) and eventually even surpasses the primary effect. By November, the ratio of switchers is 15 percentage points lower than in January, but it starts to rise again in December as the next campaign period arrives.

2.5 Discussion

An important empirical question regarding consumer attention is whether its allocation is optimal across decision problems. Using a natural experiment created by a change in auto insurance regulation in Hungary, I show that increasing the salience of a decision problem without transmitting relevant information has a large effect on people's actions.[27] Therefore, their choice to ignore the problem when it was not salient must have been strongly suboptimal.

At the same time, I also find evidence of rational elements in inattentiveness. The effect of the media campaign is proportionately stronger on those who stand to gain less from switching. Assuming that attentive people make decisions that are beneficial to them, the only explanation is that of a composition effect: low-savers were less likely to pay attention to switching in the absence of the campaign, therefore there are more low-savers among those who only pay attention because of the campaign. On balance, however, the results suggest that many people act against their best interests when they ignore the switching decision.

There are a few caveats regarding the chapter's results. First, the sample does not represent the entire population, or even the average driver. It only contains people who have acquired a car in 2010, and have used a given insurance broker to chose a liability insurance contract. It is not clear which way (if at all) the salience effect is distorted in the sample relative to the whole population. On one hand, insurance brokers tend to remind their clients of their switching deadlines, which would raise

[25] I dropped the insignificant month-benefit interaction terms from the regression.

[26] I added the control groups in pairs to save space. The excluded columns are just like the ones that are included.

[27] To prove that important contract features (most of all the prices) cannot be communicated to consumers, I collected the 2011 pricing tables of insurance companies into an online appendix, downloadable (in Hungarian) from `www.andraskiss.com/research`.

Table 2.6: Logit marginal effects on the probability of switching at the first anniversary of the insurance contract

	(1)	(2)	(3)	(4)
February	-0.056***	-0.057***	-0.056***	-0.052**
	(0.021)	(0.021)	(0.021)	(0.021)
March	-0.070***	-0.068***	-0.067***	-0.064***
	(0.019)	(0.019)	(0.019)	(0.019)
April	-0.087***	-0.087***	-0.087***	-0.085***
	(0.018)	(0.018)	(0.018)	(0.018)
May	-0.086***	-0.086***	-0.085***	-0.081***
	(0.018)	(0.018)	(0.018)	(0.018)
June	-0.106***	-0.105***	-0.105***	-0.102***
	(0.016)	(0.017)	(0.017)	(0.017)
July	-0.111***	-0.107***	-0.107***	-0.104***
	(0.016)	(0.017)	(0.016)	(0.017)
August	-0.134***	-0.133***	-0.133***	-0.132***
	(0.015)	(0.015)	(0.015)	(0.015)
September	-0.131***	-0.128***	-0.128***	-0.126***
	(0.015)	(0.015)	(0.015)	(0.015)
October	-0.140***	-0.137***	-0.137***	-0.134***
	(0.014)	(0.014)	(0.014)	(0.015)
November	-0.156***	-0.155***	-0.159***	-0.159***
	(0.013)	(0.013)	(0.013)	(0.013)
December	-0.097***	-0.093***	-0.103***	-0.100***
	(0.024)	(0.024)	(0.024)	(0.025)
Switching benefit ($100)	0.093***	0.100***	0.096***	0.078***
	(0.008)	(0.010)	(0.011)	(0.012)
Risk controls		Yes	Yes	Yes
Personal controls		Yes	Yes	Yes
Vehicle controls			Yes	Yes
Payment controls			Yes	Yes
Contract channel controls				Yes
Insurer controls				Yes
Observations	7,375	7,375	7,375	7,375

Notes: All contracts start between January 2 and December 31, 2010.

10.14754/CEU.2015.04

the non-treated switching rate and mitigate the observed effect of inattention. On the other hand, recent car buyers are younger on average, and might be more easily persuaded by the campaign to shop around.

Second, measurement error in prices could be a potential source of bias. Although I took precautions to limit the sample to those contracts for which I could calculate one of the prices accurately, I could have still made mistakes regarding the prices of the non-chosen alternatives. Most likely, calculation errors caused imputed prices to be more dispersed than they were in reality, and therefore enlarged the estimated benefits of switching. Robustness checks suggest, however, that this upward bias is limited; moreover, it does not noticeably affect the campaign estimates.

Third, I only look at switching decisions in the first year of the new regime. It is plausible that people learn from their mistakes over time and the salience effect of the campaign will vanish.[28] However, there are also several counterarguments that paint a less optimistic picture. First, many insurees are long-time participants in the system, many of whom should already be well-informed about the potential gains to switching. Yet, I find no difference in the effect of the campaign on more experienced versus less experienced drivers: both benefit from the reminders to the same extent. Second, many people have other types of insurance products (own damage and theft for vehicles, home insurance) that work with the same yearly insurance period structure and require similar 30-day notifications for cancelling, but have never been synchronized across the country. Despite the presence of insurance brokers and price comparison sites, (unconditional) switching rates in these markets are much lower than for auto liability insurance, and insurer margins are much higher. The likely implication is that people learn slowly, if at all, from errors of omission.

The results of the chapter suggest ways to think about policies to improve consumers' decision making in similar markets. Helping people judge the expected benefits better by making an (unrequested) rough price calculation for all alternatives would probably be useful, for example.[29] The coordination of all switching activity into a single month in the old regime was also a good idea, as the campaign effect estimates confirm.

Requiring people to start paying at least the first installment of next year's insurance policy *before the switching window closes* would probably also help them concentrate on the decision problem.[30] This policy suggestion is also supported by Finkelstein's (2009) study, who found less tolerance for highway toll raises whenever drivers had to pay by cash on the spot. Finally, a requirement to send a regulator-designed information leaflet on contract switching along with the announcement of next year's continuation prices could make a difference as well.

[28] Unfortunately, my dataset is too short to test for learning.

[29] Insurance brokers are usually disinclined to carry out such an exercise, for fear of being held responsible for the accuracy of the results (the rough calculation could easily be based on partially stale input data). Brokers prefer that drivers update their data first, and then show them the offers. Perhaps even more importantly, since commissions are proportional to the insurance premia, brokers are not too interested in encouraging switching as long as they can keep a client.

[30] Under current rules, the first installment is only due three months after the switching deadline.

10.14754/CEU.2015.04

2.6 Conclusions

I exploited a change in auto liability insurance regulation in Hungary to measure the causal effect of an advertising campaign on drivers' propensity to switch contracts. I showed that the campaign provides no decision-relevant information to consumers, but increases the salience of the switching opportunity, which only matters if people are suboptimally inattentive to decision making. My main result is that the media campaign has a large effect, increasing switching rates by 12 percentage points from a baseline of 20 percent. This total effect can be broken down into an equal-sized primary and a secondary element, roughly corresponding to direct exposure to media messages and message retention over time. The estimates are robust to a large variety of specifications.

The results of the chapter suggest that boundedly rational elements in decision making can have strong influence on consumer behavior, even in relatively simple settings. The main cognitive requirement in the market the paper considers is to provide basic information online for about 10-15 minutes, pick the lowest price from a list of otherwise homogenous offers, and not to miss a deadline for doing all of this. Yet, a significant proportion of people still pass up the opportunity to change contracts and leave sizeable financial gains on the table.

10.14754/CEU.2015.04

Chapter 3

Measuring Switching Costs in the Hungarian Auto Liability Insurance Market

3.1 Introduction

People often remain in high-priced contracts for insurance, utilities, or telecom services despite the availability of cheaper alternatives. This inertia is usually rationalized as the result of high costs in searching for rival offers and in administering the switching. In some contexts, however, people might simply not be aware that they have an opportunity to switch, calling into question the validity of switching cost estimates stemming from standard multinomial choice models.

In Chapter 2, I exploited a change in auto liability insurance regulation in Hungary to identify the causal effect of a media campaign that provided no decision-relevant information to consumers, but increased the salience of the switching opportunity for a well-defined time period. The campaign turned out to have a primary influence on switching, increasing switching rates by 12 percentage points from a baseline of 20 percent. I have also found that consumers were insensitive to financial incentives: the difference in switching rates between people whose savings are at the 90th and at the 10th percentile ($12 and $110, respectively) was only 8 percentage points.

In this chapter, in order to understand the effects of the media campaign, I build and estimate a structural model, in which switching costs and inattention influence switching decisions through separate channels. My estimates indicate that inattention to the switching opportunity is widespread. Without the campaign, two-thirds of consumers ignore the decision problem altogether, whereas during the campaign the implied ratio of inattentive people is less than one half. Estimated mean switching costs are around $65, a plausible number given industry reports. However, the failure to account for the presence of inattentive consumers biases switching cost estimates upwards by an order of magnitude, which explains the

10.14754/CEU.2015.04

apparent insensitivity to financial incentives in the reduced-form specifications.[1]

I incorporate the idea of random attentiveness into a structural switching cost model, and show that the augmented framework greatly improves the plausibility of the estimation results and yields new insights into consumers' behavior. Specifically, I construct a two-period random utility model in which people make standard multinomial choices among insurance companies. In the first period, the decision problem is symmetric, as there are no default options. In the second period, consumers must pay an extra cost if they want to switch to a different insurance contract. I also assume that the second period choice is taken only with probability θ (the "attention parameter"). With probability $1-\theta$, people remain inattentive to the decision problem, and the default contract continues automatically for another year. I allow both the switching cost and the attention parameter to depend on individual characteristics, and estimate the augmented choice model using maximum likelihood methods.

The effects of switching costs and inattention on choices are identified from the way they influence people's responsiveness to financial savings. When switching costs rise, they tend to affect people with intermediate savings the most. Those with low savings will rarely switch, and those with high savings will always switch, so the response comes from those whose savings are on the margin. On the other hand, inattentiveness—by definition—is equally likely regardless of financial savings, and hence it elicits a stronger switching response from high savers.

The main result of the structural estimation is that the switching campaign mainly acts through manipulating attention, rather than switching costs. Only one in three people consider switching without the campaign, whereas more than half of them do so during the campaign. Switching costs, on the other hand, are at around \$65 in both periods. In contrast, personal characteristics, such as age or vehicle power tend to influence switching costs, rather than attention levels.[2]

A handful of recent empirical papers have tried to separately measure the sources of consumer inertia. In contemporaneous work, Hortaçsu et al. (2014) look at switching decisions in the Texas residential electricity market using monthly consumption data for households. They specify a two-stage discrete choice model to separate inattention from brand preferences, and conclude that people only pay attention to the supplier switching decision once every 4-5 years. My structural model setup closely resembles theirs and the results I get have similar magnitudes, even though the data, the sources of identification, and the estimation methods are different in the two studies. By using changes in market regulation as a natural experiment, however, I am also able to measure how much a real policy change can influence consumer decisions by decreasing the share of inattentive people.

In a similar context, Miravete and Palacios-Huerta (2014) attempt to separate

[1] A note on terminology: in a narrow sense, switching costs only refer to the cost of changing service providers once all the offers are known, search costs refer to the cost of learning about the alternatives, whereas the value of consumer inertia is an encompassing term for all barriers to switching. Theoretical switching cost models, on the other hand, tend to use the term "switching costs" to mean consumer inertia in general, mostly because they abstract away from the search process. In this chapter, switching costs are also used as a remainder quantity: once inattention and persistent preferences for insurers are taken into account, switching costs represent the financial value of everything else that prevent people from switching, including the cost of obtaining offers.

[2] Age increases switching costs by \$12 for each decade. Having a more powerful car lowers switching costs by \$25 for each standard deviation in power.

10.14754/CEU.2015.04

the effect of endogenous past experience and learning from the effect of pure inertia on tariff plan switching decisions in a local telephone market in the U.S. Their overall finding is that households tend to learn from their mistakes and make better decisions over time, and those who face cognitively less demanding choice problems learn faster. The authors argue that while inertia still exists among consumers, it is likely caused by rational inattention to the switching problem. In contrast, I show that people are more likely to make good financial decisions when the choice situation is more salient, and hence it is *not* rational to be inattentive for a large share of the population.[3]

The method I use to separately identify inattention from financial costs of switching may be more generally applicable to other choice situations. There are a large number of studies that structurally estimate fixed costs associated with stickiness in labor economics (Artuç et al. (2010)), international trade (Das et al. (2007)), or monetary macroeconomics (Golosov and Lucas (2007)). Generally, these studies estimate large fixed costs that prevent agents from choosing better alternatives. It may instead be that true fixed costs are lower, but the presence of agents who suboptimally ignore the decision problem biases the fixed cost estimates upwards. Using the structural approach of this chapter, the bias could be corrected.

The rest of the chapter is organized as follows. Section 2 provides a background on the vehicle liability insurance market in Hungary. In section 3, I describe the data and the identification of the campaign effect, and substantiate the comparability of people in the distinct treatment groups. Both of these sections are intentionally short; the same ideas are explored in greater depth in Chapter 2. Section 4 contains the main contribution of the chapter, a structural switching cost model with inattentive consumers. I report the main results and the most informative robustness checks in Section 5, followed by an evaluation of counterfactual policy experiments, and concluding remarks.

3.2 Vehicle liability insurance in Hungary

In Chapter 2, I provide a detailed description of the Hungarian vehicle liability insurance market. The current section is an abbreviated refresher of the rules and regulations by which the market operates.

Auto liability insurance in Hungary is an insurance product standardized by law. It covers property damage and bodily injuy to third parties in an accident. Coverage to the faulty driver's car and person must be purchased separately, if desired.

Vehicles must be insured to be allowed on the road. Companies are free to set their own prices (with certain restrictions on the price structure and the admissible pricing factors), and there are over a dozen insurance providers on the market. Price setting rules are public and (during the period of my study) announced simultaneously for an entire calendar year by the insurers.

[3]Luco (2014) and Honka (2014) also separate certain elements of consumer inertia using data from pension funds in Chile and auto insurance in the U.S. Luco (2014) finds that the costs evaluating financial information and making decisions is twice as large as the administrative costs of switching between fund managers. Honka (2014) jointly estimates a search and switching model and reports that search costs are the main determinants of consumer retention in the U.S. auto insurance industry.

10.14754/CEU.2015.04

There are two main reasons why a new insurance contract is started: (1) the acquisition of a new vehicle, and (2) contract switching on an anniversary. Contracts have indefinite length by default, but when a vehicle is sold, the contract ends automatically. Otherwise, people must wait until the next anniversary of the initial contracting date to terminate the contract.

Insurees who want to switch providers must provide a 30 days' cancellation notice before the anniversary, otherwise the contract continues automatically for another year. Drivers can learn about the continuation price and all the alternative offers at least 60 days before the contract turns for another year. If the contract is cancelled in time, people are free to choose any another contract—including the one offered to new customers by their previous insurance provider.

Searching for alternatives can be prohibitively costly if done "by hand", but there are many insurance brokers who specialize in automating the price comparison process. Insurance brokers are often independent of insurance companies, take a cut from the final price (paid by the insurers), and do not ask for client payments. Their commissioning is therefore free of charge to drivers, and it can substantially speed up both the choice among the alternatives and the subsequent contracting administration.

Prior to 2010, contracts were required to coincide with the calendar year from the second year onwards, meaning that the time window for contract switching was synchronized for all drivers to November. This particular regulation fuelled a concentrated switching campaign by insurance companies and brokers. Following January 1st, 2010, the synchronization is not carried out any longer, and people must adhere to the anniversary of their initial contract start date with their switching decisions.

3.3 Data

I use the database of a mid-sized independent insurance broker company in Hungary. The company keeps a record of all personal and vehicle characteristics relevant to a liability insurance contract, as well as general contract features.[4]

I restricted the original multi-year database to contracts that started in 2010 as a result of a vehicle acquisition. This is important, because it lets me observe an initial choice without a default option, and hence without switching costs.

Moreover, the choices in 2010 are uncensored, whereas in 2011, I mostly do not observe the identity of the company that switchers have chosen (about 90 percent of the 2011 data on switchers is censored). Consequently, the estimates for the insurer fixed effects and the price sensitivity of drivers are mainly driven by the first period (2010) choice, whereas the second period (2011) choices pin down the parameters that govern the switching decision.[5]

Regarding contract price data, my dataset is limited. I only see the price of the chosen contract in 2010, but not the continuation price in 2011, or the alternative

[4]Unfortunately, I only have access to a subset of all relevant characteristics, which introduces measurement error into the calculation of unobserved alternative prices.

[5]Handel (2013) makes a similar point about the importance of observing initial choices without default options in the estimation of subsequent consumer inertia.

10.14754/CEU.2015.04

Table 3.1: Descriptive statistics

	Count	Mean	St.dev.	Min.	Max.
Age	6,766	40.26	12.93	16.00	91.00
Female	6,766	0.326			
Capital resident	6,766	0.136			
Years since license obtained	6,267	17.79	11.29	0.00	58.00
A0	6,766	0.773			
B10	6,766	0.107			
Cylinder capacity (ccm)	6,766	1,552	412	594	4,973
Vehicle age (years)	6,766	9.441	5.309	0.000	48.000
Car dealership contracts	6,766	0.926			
Online contracts	6,766	0.044			
Payment by check	6,766	0.937			
Quarterly payment	6,766	0.940			
Insurance premium ($100)	6,766	1.859	0.926	0.381	9.945
Switching benefit ($100)	6,766	0.567	0.497	-0.169	5.735
Switching (Jan 2010 contracts)	326	0.374			
Switching (Feb-Nov 2010 contracts)	6,440	0.203			

offers in either 2010 or 2011. Fortunately, the pricing tables of the insurance companies can be used together with the observed personal and vehicle characteristics to reconstruct all offers in both years—with a certain measurement error. To limit the influence of this error on the estimates, I restrict the final sample to only include those contracts for which the single verifiable reconstructed price (the chosen offer in 2010) is reasonably accurate.[6]

Descriptive statistics about the estimation sample are shown in Table 3.1.[7] For binary variables, I only report the sample means (i.e. the share of people with the given characteristics in the sample). For all other variables, the table also contains standard deviations, as well as minimum and maximum values.

3.4 Structural model for insurance switching

The reduced-form results established that the salience of the switching opportunity matters for switching decisions. What is still unclear, however, is whether people act as fully attentive consumers during the campaign period, or many of them remain

[6] One potential worry in the estimation is the endogeneity of prices. In a usual demand analysis, price endogeneity is a concern when aggregate data are used for the estimation, but this is not the case here, since I observe individual decisions. However, the endogeneity issue does arise in a different form. Price setting is based on people's observable characteristics, but some of these characteristics are missing from my dataset. An unobserved price-proportional discount, for example, introduces positive correlation between the calculated price and the error term of the regression. My main defense against such price endogeneity is the sample restriction to limit the measurement error in prices. Robustness checks using a range of error tolerances suggest that price endogeneity is not a serious concern in my dataset.

[7] The appendix of Chapter 2 contains more information about sample construction.

Figure 3.1: Switching rates as a function of financial savings

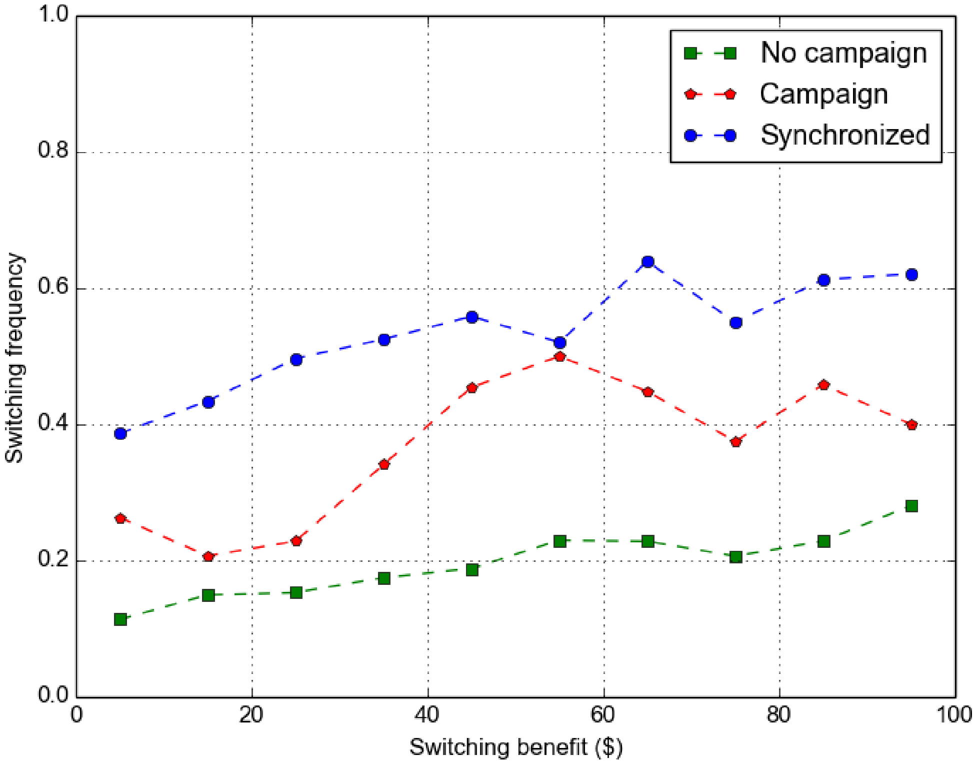

ignorant of switching despite the campaign. In this part of the chapter, I build and estimate a structural choice model to answer this question by separating the effects of switching costs and inattention on consumer inertia. I show that this augmented framework greatly improves the plausibility of the estimation results and yields new insights into consumers' behavior.

Figure 3.1 motivates the need to build inattention into a standard choice model. The horizontal axis in the figure shows the financial savings from switching. On the vertical axis, I plot the average switching rates on new 2010 contracts at the end of the first insurance period. The averages are calculated by grouping the savings variable into $10-wide bins.[8] The three separate curves—from bottom to top—show (1) new-regime drivers who are not directly affected by the campaign, (2) new-regime drivers who are directly affected by the campaign, and (3) old-regime drivers who cancelled their previous contracts in November 2009 and signed a new one on 1/1/2010 (the *Calendar* group).

As we would expect, switching rates rise with the amount of savings: people do respond to financial incentives. Moreover, when the switching opportunity is more salient, a larger share of the population decides to switch for the same level of savings, which is in line with the reduced-form results. But Figure 3.1 also highlights

[8]There is one exception: the rightmost point on the graph shows the average switching rates of all people whose savings exceed $90. The number of observations drop sharply beyond $100.

10.14754/CEU.2015.04

another phenomenon: many people seem to be unwilling to switch insurance contracts regardless of how much they can save by doing so. Graphically, this shows up in the "flattening" of the switching rate curves as the level of savings goes above approximately \$60. The switching rates of even the most active consumers, the old-regime drivers, top out at about 60 percent.

It seems implausible that these people are fully aware of the potential savings and the effort involved in switching, and have consciously decided that they would rather forgo \$100-150 than attend to the decision problem. More likely, they do not pay attention to the decision at all, regardless of how much they could save by acting. This ignorance also explains why financial incentives lose their power well before switching rates approach 100 percent.

For the purpose of structural modeling, I define attentiveness to a decision problem as the conscious utility maximizing choice among all available options. Conversely, inattentiveness means that the decision maker will not examine any of the available alternatives, and therefore his existing contract will continue for another insurance period by default.[9] I assume that attention is a random variable, which may depend on the characteristics of individuals and of the environment in general, but is independent of the characteristics of alternatives. In particular, attention does not depend on how much a person could save by being attentive.

Specifically, I construct a two-period random utility model in which people make multinomial choices among insurance companies. In the first period, the decision problem is symmetric, as there are no default options. In the second period, consumers must pay an extra cost if they want to switch to a different insurance contract. In addition, the second period choice is taken only with probability θ (the "attention parameter"). With probability $1-\theta$, the default contract continues automatically for another year. I allow both the switching cost and the attention parameter to depend on individual characteristics. I estimate the structural model using maximum likelihood methods.

3.4.1 Model

Initial choice for all consumers

In the initial period (denoted by $t = 0$), drivers choose an insurance contract without having a default "do-nothing" option. This corresponds to new-regime insurees buying a car in 2010, for which liability insurance must be acquired from the day of the purchase.[10]

Specifically, I assume that driver n receives the following (indirect) utility by choosing contract j at time 0:

$$U_{nj0} = \alpha X_{nj} - p_{nj0} + \nu_0 \varepsilon_{nj0} \tag{3.1}$$

p_{nj0} is the insurance premium that driver n would pay in contract j in 2010, and it affects utility negatively. X_{nj} are insurer dummy variables; hence α captures the

[9]Of course, an attentive person can also come to the conclusion that the benefits of switching are not worth their costs, and hence remain with the default alternative.

[10]The "no-default" modeling assumption in the initial period means that I do not seek to fit the behavior of old-regime drivers in Figure 3.1 to the model.

insurer fixed effects. In the baseline specifications, $X_{nj} = X_j$, meaning that insurer fixed effects are constant across people. I allow for individual specific fixed effects during the robustness checks, but restrict the distribution of X_{nj} to be normal with j-specific means and variances. Individual-specific fixed effects control for unobserved preference heterogeneity for insurers, but it turns out that the $X_{nj} = X_j$ restriction does not make noticeable difference to the results.[11] The coefficient of the price variable is set to 1, normalizing the measurement unit of all other coefficients to hundred-dollar terms.

Utility also contains a random term ε_{nj0}. For tractability, I restrict ε_{nj0} to be independently and identically distributed across drivers and insurers following a Type-I extreme value distribution. Since the scale of utility is set by the price normalization, the variability of the random term is an estimable parameter, denoted by ν_0. The setup in period 0, therefore, is that of a typical conditional logit model as described by Train (2003), among others.

The probability of person n choosing contract $j \in J^0$ at $t = 0$ is:

$$P_{nj0} = \frac{\exp\left(\frac{\alpha}{\nu_0} X_{nj} - \frac{1}{\nu_0} p_{nj0}\right)}{\sum_{i \in J^0} \exp\left(\frac{\alpha}{\nu_0} X_{ni0} - \frac{1}{\nu_0} p_{ni0}\right)} \tag{3.2}$$

where J^0 denotes the set of all available contracting alternatives in the initial period.

Switching decision for attentive consumers

In period 1 (corresponding to 2011), drivers have an option to stick to their previously chosen contract, or they can cancel the contract and sign a new one with any of the insurance companies, including their previous provider.

Utility in period 1 is specified as:

$$U_{nj1} = \alpha X_{nj} - p_{nj1} - \beta Z_{nj1} + \nu_1 \varepsilon_{nj1} \tag{3.3}$$

The insurer fixed effects are time-invariant.[12] The price variable p_{nj1} reflects what the consumer would have to pay in period 1 for contract j, while the random term ε_{nj1} has the same EV Type-I distribution as its—independent—period 0 counterpart. The variability of the random utility component is allowed to be different in the two time periods.[13]

Z_{nj1} contains the determinants of switching costs, interacted with a switching indicator. In the simplest case, $Z_{nj1} = 1$, and the corresponding β parameter (the

[11] On the other hand, the restriction allows for much easier estimation, since it does not require simulation-based techniques.

[12] For the sake of simplicity, there is a slight abuse of notation at this point, because I do not explicitly distinguish between contracts and insurers. The one-insurer-one-contract equivalence breaks down when a driver stays at the same insurer, but signs a new contract. This is treated as switching, but the insurer fixed effect is the same in both cases. Hence the number of options available to drivers equals the number of insurers plus one, whereas the number of insurer fixed effects to estimate only equals the number of insurers minus one (one of the fixed effects must be normalized).

[13] In an earlier version of the paper, I restricted the variance of the random utility component to be uniform over time. It turns out that this restriction heavily distorts the switching cost estimates towards zero. I will discuss this issue in the section on robustness.

intercept) measures the average costs of switching across the sample. The period 1 choice probabilities are analogous to expression (3.2).

The independence assumptions on the random utility components (ε_{njt}) are not entirely inocuous. Independence across alternatives (j) might fail, for example, if a person was eligible for unobserved discounts offered by some insurers, but not by others. Since the discount structures of companies are highly persistent, independence over time might fail as well in such a scenario. I have a two-fold strategy for mitigating the problem. First, I restrict the sample to minimize measurement error in prices, which is equivalent to minimizing unobserved discounts. Second, I allow for insurer fixed effects to differ across people (although they remain uncorrelated across alternatives). My estimates are robust to a range of these error-mitigating methods, implying that the independence assumptions on error terms are not unreasonable.

There is also a second issue related to choice over time. In my model, people are myopic: they optimize in each year, but not for both years at the same time. I view this assumption to be by far the most realistic for the market. Companies set prices for hundreds of market segments (based on age, car type, location, driving history, etc.) simultaneously, once a year. Costs (accident damages) are random and change over time even in expectation (safer driving, better roads). Prices are widely dispersed and there is no obvious low-price firm in the market, implying that firms are unlikely to play pure strategy equilibria. Under such conditions, and considering that switching is relatively easy and costless, forming expectations about future prices is a pointless exercise for even the most sophisticated consumers.

Inattention to switching

In period 1, I allow consumers to be randomly attentive to the switching decision in the following way:

$$j_n^1 = \begin{cases} \underset{j}{\operatorname{argmax}}\, U_{nj1} & \text{with prob. } \theta_n \\ j_n^0 & \text{with prob. } 1-\theta_n \end{cases} \tag{3.4}$$

where j_n^t is the chosen alternative by consumer n at time t, and the probabilistic attention parameter θ_n is defined by the following functional form:

$$\theta_n = \frac{1}{1+\exp(\gamma W_n)} \tag{3.5}$$

Thus, a consumer takes a conscious utility-maximizing choice (allowing for switching costs) in period 1 with probability θ_n, and simply continues his existing contract with probability $1-\theta_n$. The attention parameter may depend on individual characteristics (W_n), such as whether a person was affected by the campaign or not. In the simplest case, when W_n is only a constant, $\frac{1}{1+e^{\gamma}}$ measures the share of people who consciously consider the switching decision in the entire sample.

Structural identification

The effects of switching costs and inattention on choices are identified from the way they influence people's responsiveness to financial savings. Figure 3.2 provides a

10.14754/CEU.2015.04

stylized graphical explanation.

The horizontal and vertical axes in Figure 3.2 measure the financial savings and the probability of switching, similarly to Figure 3.1. The logistic-shaped curves show the reaction to financial savings in a simplified setting, assuming only two alternatives. The estimated costs of switching are denoted by c.

The top panel of Figure 3.2 shows a scenario where inattentive people are not present and the campaign has its effect on observed switching rates by decreasing the costs of switching. The change in switching costs mostly affects the switching decisions of people with intermediate savings. Those with low savings rarely switch, and those with high savings always switch, so the response comes from those whose savings are on the margin of surpassing the costs of switching. In geometric terms, changes in switching costs stretch or compress the switching reaction curve in a horizontal direction, but do not affect its lower or upper limits.

In the bottom panel of the figure, not all insurees are attentive, hence the upper limits of the switching reaction curves are below 100 percent. Moreover, the campaign works through increasing the probability of attentiveness, measured by the increase in the limit to which the switching reaction curve converges as financial savings grow. Geometrically, changes in attention stretch or compress the switching reaction curve in a vertical direction, and therefore the percentage point increase in switching rates would be largest for people with high savings.

Intuitively, the identification of the attention parameter comes from the upper limit that is fitted to the curves in Figures 3.1 and 3.2, whereas the switching cost equals the amount of financial savings for which the observed switching rates are equal to those predicted by the model.

Censoring

There is one additional complication in the estimation procedure, which results from limitations in the data. Although I observe whether people switch insurance contracts in 2011, for the majority ($\sim$ 90%) of switchers I do not know which company they have switched to. This censoring is a result of how the data is generated,[14] and it can reasonably be assumed to be independent of the (unobserved) choices themselves. As a result, the fact that an observation is censored yields no additional information about which non-default alternative was chosen, compared to what I already observe about the alternatives.

The censoring slightly modifies the choice probability for a switching decision, because I have to add up the probabilities of all the unobserved choices. For a censored observation, I can only apply one of two choice probabilities (switching or not), whereas for a non-censored observation, I can pick from as many probabilities as there are alternatives to choose from. Censoring is not an issue in the period 0 decision.

[14]People are free to choose among insurance brokers to conduct the administrative process of switching on their behalf. When a person uses a different broker, I know that he is switching contracts, but the new contract (and hence the identity of the insurer) will not be in my dataset.

Figure 3.2: Structural effects on switching costs vs attention

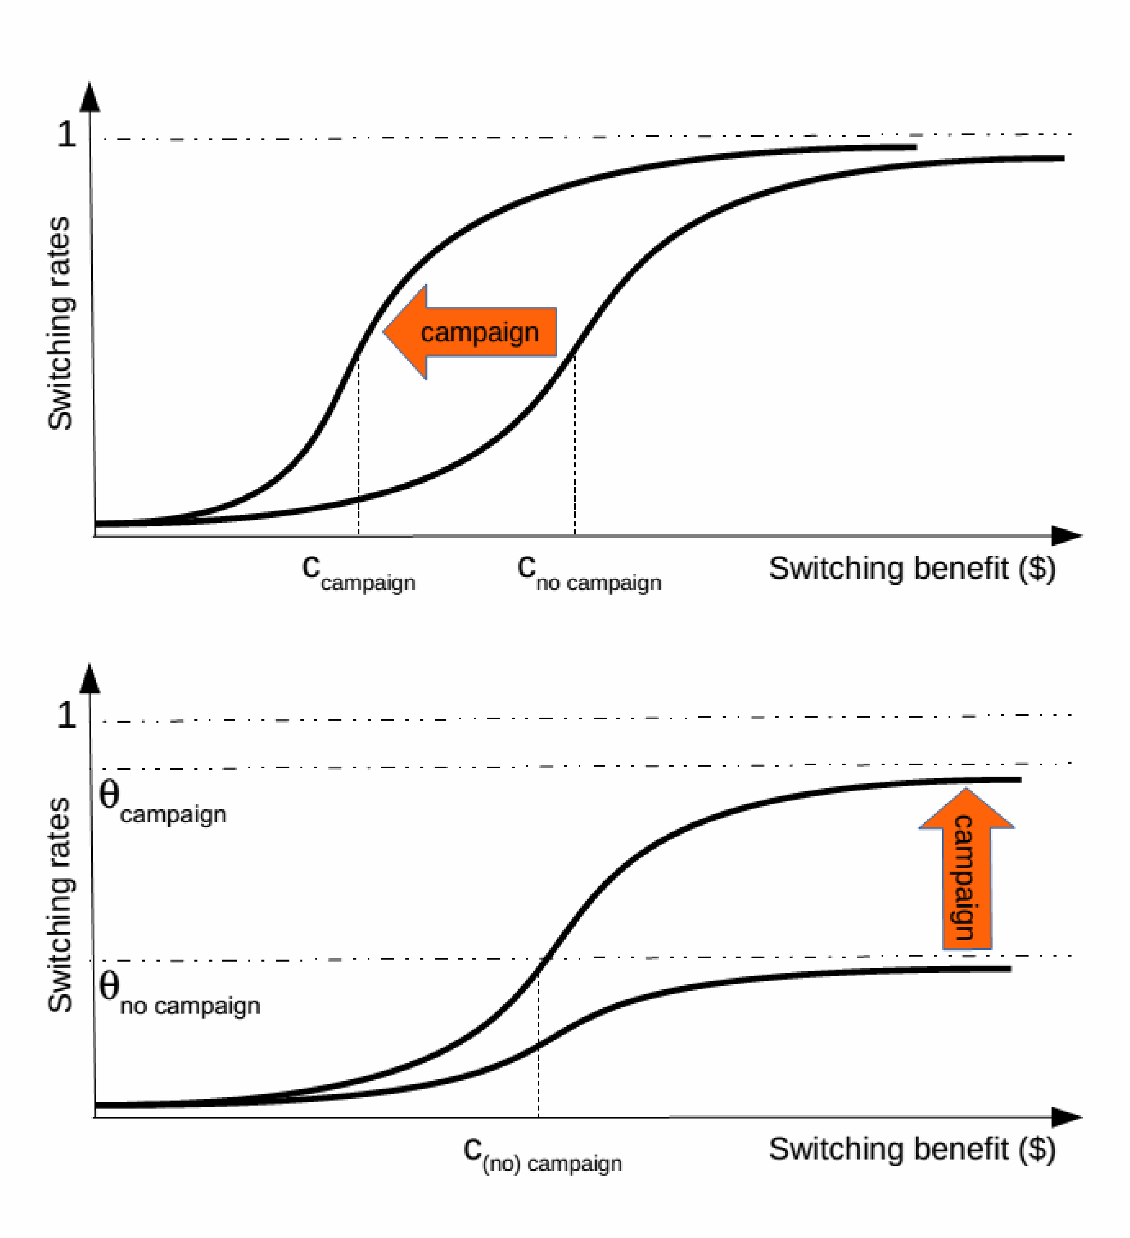

10.14754/CEU.2015.04

Aggregate choice probabilities

Putting together the observations from the two periods, we get the following expression for the aggregate choice probability:

$$P_n = P_{nj^0 0} \cdot P_{nj^1 1} \tag{3.6}$$

where $P_{nj^0 0}$ is given by equation (3.2) for the actual choice $j = j^0$, and

$$P_{nj^1 1} = \theta_n \cdot \frac{\exp\left(\frac{\alpha}{\nu_1} X_{nj^1} - \frac{1}{\nu_1} p_{nj^1 1} - \frac{\beta}{\nu_1} Z_{nj^1 1}\right)}{\sum_{i \in J^1} \exp\left(\frac{\alpha}{\nu_1} X_{ni} - \frac{1}{\nu_1} p_{ni1} - \frac{\beta}{\nu_1} Z_{ni1}\right)} + (1 - \theta_n) \cdot d_n \tag{3.7}$$

when an observation in period 1 is not censored, and

$$P_{nj^1 1} = \theta_n \cdot \left[1 - \frac{\exp\left(\frac{\alpha}{\nu_1} X_{nj^0} - \frac{1}{\nu_1} p_{nj^0 1} - \frac{\beta}{\nu_1} Z_{nj^0 1}\right)}{\sum_{i \in J^1} \exp\left(\frac{\alpha}{\nu_1} X_{ni} - \frac{1}{\nu_1} p_{ni1} - \frac{\beta}{\nu_1} Z_{ni1}\right)}\right] \tag{3.8}$$

when an observation is censored. In the second case, all we know is that the original contract was *not* chosen. In both equations, d_n denotes whether a person has chosen his default option, or not ($d_n = 1$ for non-switchers, and 0 for switchers). Censoring automatically implies that a person has switched contracts.

3.4.2 Estimation

Taking the logarithm of expression (3.6) and summing over the individuals yields the objective function of a maximum likelihood estimation procedure.

$$ll = \sum_n \log\left(P_{nj^0 0} \cdot P_{nj^1 1}\right) \tag{3.9}$$

When the structural model has deterministic coefficients only, the log-likelihood function has a closed form and can be maximized using standard optimization procedures.[15]

On the other hand, when we allow some of the structural parameters to be drawn randomly for each person (but from distributions with known shapes), the objective function no longer exists in closed form. Instead, we have to resort to simulated maximum likelihood by maximizing the following expression:

$$sll = \sum_n \log\left(\frac{1}{R}\sum_{r=1}^{R} P^r_{nj^0 0} \cdot P^r_{nj^1 1}\right) \tag{3.10}$$

[15] However, during exploratory runs on artificial data I have found that several of these standard procedures (namely, the *fminunc* function in Matlab, and the *fmin_bfgs* function in Numpy/Python) ran into difficulties in the optimization process and could not find a solution. This was most likely a consequence of non-concavities in the objective function further away from the optimum. To get around the problem, I developed a custom optimization routine, which uses a combination of the Newton-Raphson and the BHHH (Berndt et al. (1974)) procedure with automatic step-size adjustment. The routine performs well on both artificial and real data, finding the same solutions from a variety of starting points. The code is programmed in Python and is available upon request.

where $P^r_{nj^0j}$ and $P^r_{nj^11}$ are the probabilities that person n makes the choices we observe, given that the random parameters take the values that resulted on the rth draw. Expression (3.10) is the sample counterpart of "integrating out" the random parameters from the log-likelihood expression (see Ch.6 in Train (2003) on mixed logit models).

More specifically, let us partition all the individual parameters within α, β, γ, and ν into two groups: those that are deterministic (δ) and those that are allowed to vary randomly across people (ρ). In each iteration of the simulated maximum likelihood procedure, I calculate the value of *sll* the following way. I take R draws from the joint distribution of ρ, which I assume to be multivariate normal with unknown means, unknown variances, and zero correlation across any two elements.[16] For each draw (ρ^r), I substitute δ and ρ^r into equation (3.6) to get a joint choice probability for person n in the two periods. For each person, I average these joint choice probabilities across the R draws, which will thus depend on the deterministic parameter vector δ and the means and variances of the random parameter vector ρ. I take the logarithm of the average simulated choice probabilities and sum them over the entire sample to arrive at the simulated log-likelihood value that is to be maximized.

Although tested routines for mixed logit estimation on panel data are publicly available,[17] I had to develop my own version for two reasons. First, the existing codes (that I am aware of) do not deal with the special case when the number of alternatives varies by years and individuals.[18] Second, I had to account for the censored observations properly in the simulated log-likelihood function.[19]

3.5 Results

3.5.1 Baseline estimates

Table 3.2 contains the main results of the structural estimation. The regression table is composed of several panels, which are separated by dashed horizontal lines. The columns show the results for different sets of explanatory variables included in the switching cost specification.

The top panel contains the most interesting switching cost effects (denoted by β in equation 3.3 and measured in \$100), the second panel shows the estimated parameters influencing inattention (γ), the third panel marks the applied sets of control variables, followed by the estimates of the random scale parameters (ν). The number of observations is the same as in the main reduced-form specifications.

The intercept of the switching cost component shows that estimated switching costs are around \$65, which is a plausible outcome given industry reports in trade

[16]In practice, $R = 200$ and I use pseudo-random numbers to generate the normal distributions. Trial runs on artificial data of similar magnitude suggest that this method produces sufficiently accurate estimates.

[17]See, for example, the Matlab codes on Kenneth Train's website: `http://eml.berkeley.edu/Software/abstracts/train1006mxlmsl.html`.

[18]In my data, there are more insurers in period 1, than in period 0, as well as an additional default alternative. Moreover, one insurance company could not be chosen in the first 3.5 months of 2010 for administrative restrictions.

[19]The mixed logit codes were programmed in Python, and are available upon request.

10.14754/CEU.2015.04

Table 3.2: Parameter estimates in the structural switching cost model with inattentive consumers

	(1)	(2)	(3)	(4)
Intercept (switching cost)	0.653***	0.695***	0.627***	0.646***
	(0.111)	(0.116)	(0.121)	(0.122)
Age		0.013***	0.013***	0.012***
		(0.003)	(0.003)	(0.003)
Online contracts				0.330
				(0.234)
Campaign	-0.102	-0.105	-0.068	-0.033
	(0.250)	(0.256)	(0.240)	(0.227)
Power (kW)			-0.010***	-0.010***
			(0.003)	(0.003)
Intercept (θ)	0.967***	0.940***	0.986***	0.905***
	(0.082)	(0.084)	(0.078)	(0.096)
Online contracts (θ)	-0.581***	-0.596***	-0.644***	-0.876***
	(0.152)	(0.154)	(0.155)	(0.273)
Campaign (θ)	-0.905***	-0.906***	-0.917***	-0.975***
	(0.261)	(0.271)	(0.241)	(0.250)
ν_0	0.290***	0.290***	0.291***	0.291***
	(0.004)	(0.004)	(0.004)	(0.004)
ν_1	0.503***	0.516***	0.525***	0.526***
	(0.031)	(0.033)	(0.034)	(0.034)
Insurer fixed effects	Yes	Yes	Yes	Yes
Risk and personal controls				Yes
Vehicle and contract controls				Yes
Personal controls (θ)	Yes	Yes	Yes	Yes
Observations	6,766	6,766	6,766	6,766

Table 3.3: Total effects on the attention parameter of the structural model

Intercept	x	x	x	x
Campaign		x		x
Online contract			x	x
Attention probability	29%	52%	49%	72%

press articles.[20] Being 10 years older is associated with a switching cost increase of \$12, while having a car that is one standard deviation (25kW) more powerful decreases switching costs by \$25. Being affected by the campaign is immaterial for the switching decision, provided that a consumer is attentive.

The two main shifters of attention probabilities are the campaign and the contracting channel through which the initial contract was signed. Table 3.3 shows the values of θ for various combinations of attention determinants.

Baseline inattention is high: about 70 percent of insurees who sign their original contracts in a car dealership would not consider switching one year later. The campaign "treatment" persuades a third of these inattentive people (or 23% of the treatment group) to pay attention to the switching opportunity. Slightly more than half of those who are persuaded (12% percent of the treatment group) do eventually switch contracts. People who dealt with the original contract choice by themselves using the online interface of the insurance broker are 20 percentage points more attentive to switching after one year. However, even they benefit from a reminder, as evidenced by the comparison of the last two columns in Table 3.3.

These numbers line up well with the figures cited in the Scale Research (2010) study on old-regime drivers. In that survey, 28 percent of the respondents said that they would have paid attention to the switching decision with or without the campaign (cf. column 1 of Table 3.3). Another 12 percent said that they took care of switching only because of the constant reminders, whereas 17 percent of people heard about the switching opportunity, but didn't do anything. The sum of these latter two numbers could be compared to the 23% of people who think about switching as a result of the campaign in the structural model. Finally, the rest (43 percent) of the respondents neither heard, nor cared about switching insurance contracts. They are the complements to a weighted average of columns 2 and 4 in Table 3.3, where the weights correspond to the share of people using online contracting channels in the entire population.

[20]Other empirical auto insurance studies in different countries and time periods have found search and switching costs of varying magnitudes. Honka (2014) finds the median value of consumer inertia to be around \$400, although she attributes the larger part of this amount to consumer satisfaction with the current provider and to the costliness of search, and only about \$42 to the actual hassle of switching. Berger et al. (1989) have estimated insurance switching costs in the \$185–381 range, whereas Dahlby and West (1986) have found search costs between \$131 and \$570 (both sources were converted to current U.S. dollars by Honka (2014)). My estimates are generally lower than these earlier results when I take the possibility of inattentiveness into account. On the other hand, Cummins et al. (1974) have calculated switching costs to be 20 percent of the insurance premium, which is only about \$32 in my case.

10.14754/CEU.2015.04

3.5.2 Robustness checks

Heterogenous insurer fixed effects

In the main structural specification, insurer fixed effects are constant across people. However, this assumption may be too restrictive. For example, one person might have a home insurance at Allianz, and another at Generali. Since insurers give discounts for cross-sales (a form of bundling that I do not observe in the data), the first person would favor Allianz over Generali, and vice versa. Given that I do not allow insurer fixed effects to vary across individuals, I might mistakenly think that consumer inertia is due to switching costs, whereas it is a result of unobserved, but persistent, preference heterogeneity.

As a robustness check for heterogenous insurer fixed effects, I estimate the structural model using a mixed logit specification, allowing all of the α parameters to take on individual-specific (but time-invariant) values from independent normal distributions.[21] The results are shown in Table C.1 in the appendix. The estimates with heterogenous fixed effects are very similar to the baseline specification. The baseline attention rate is slightly higher (by about 3 percentage points), and the point estimates for the switching cost are lower by about \$5 (a non-significant difference). The effect of the campaign on attention is unchanged. The conclusion is that unobserved preference heterogeneity does not noticeably bias the baseline structural estimates.

No insurer fixed effects

I also estimate the model by excluding all insurer fixed effects in both periods. This setup corresponds to a situation in which—conditional on prices—people display no tendency to favor one insurance company over another (either over time, or across individuals in the same period). That is, I assume that people only care about the insurance premia and that premia are calculated without systematic errors. Since product quality is regulated, not distinguishing between insurers is actually a sensible approach to shopping in this market.

I find that the estimates for the inattention parameters are the same with or without insurer fixed effects, but the switching cost estimates rise from \$65 to about \$100 when fixed effects are excluded. Consequently, about a third of the estimated switching costs in this alternative specification are due to population-level preferences for certain insurers.

Implicit default options at the initial choice for long-term insurees

In the model setup, I argued that the initial choice (in 2010) is without an explicit default option, meaning that people have to make an active choice between all the insurance companies in the market. At the same time, drivers may have an implicit default option if they have been insured by the same company for a long time, and especially if they have coexisting insurance products with an insurer. These implicit

[21] That is, instead of estimating the α coefficients, I estimate the mean and the variance of their (independent normal) distributions.

10.14754/CEU.2015.04

default effects act like unobserved preferences for the chosen insurers, which—if disregarded—can bias switching cost estimates upwards.

Allowing for heterogenous insurer fixed effects was one way of checking whether implicit default options are important. Another robustness check is based on implicit default options being more important for long-time insurees, since they are more likely to have built up a history with specific insurance companies, and also more likely to have coexisting insurance products, such as home insurance. On the other hand, recent entrants to the auto liability insurance market are probably less affected by these considerations, and for them, the initial choice is truly without a default option.

I re-estimated the baseline model for the following two subsamples to check for differences in switching costs and attention levels: (1) people of age 36 or younger, and (2) people with a risk rating of A0, B1, or B2. Both categorizations are imperfect measures of being a recent entrant to the market.

The results of the estimation are collected in Table C.3 in the appendix. Column 1 of the table shows the baseline estimates from Table 3.2 with all controls included, whereas columns 2 and 3 show the alternative subsamples with recent entrants (people under 36 and people with low risk ratings, respectively). The results are inconclusive: estimates on the attention parameter are fairly robust across the different specifications, but the switching cost intercept does vary considerably. Although recent entrants seem to have lower switching costs, giving support to the implicit default options theory, the estimates are also very imprecise.

Full attention

We can confirm the benefit of explicitly modeling inattention by running an estimation in which θ is set to 1; that is, all consumers are assumed to be fully attentive. The result of this exercise is shown in Table 3.4. Switching costs increase from \$70 to about \$550, which is an unreasonably high value.[22] The campaign decreases these costs by about \$120, but the remainder is still substantial. In addition, the variance of the period 1 error term almost triples, signalling that the failure to account for inattention considerably worsens the explanatory power of the model.[23]

Measurement error in insurance premia

I conduct a final robustness check to verify the sensitivity of the estimates to measurement error in the unobserved alternative insurance premia. Using the public price schedules of insurance companies, the available information on insurees, and

[22] Allowing for switching costs to be individually drawn from a normal distribution, I also find that switching costs are widely dispersed: the 10th percentile is about \$20, while the 90th percentile is \$940 (the mean in the mixed logit estimation is \$480). The large dispersion is essential for capturing the presence of attentive consumers with reasonably low switching costs and of inattentive consumers with seemingly infinite switching costs in the same framework.

[23] When the variance of the random utility term $\nu \cdot \varepsilon$ is restricted to be uniform across the two time periods, while assuming that $\theta = 1$, the switching cost estimates are considerably lower: only at around \$160, instead of \$550 (see Table C.4 in the appendix). The uniform variance restriction cuts the standard deviation of the period 1 error term by over two-thirds, decreasing the dispersion in the utilities offered by alternative choices. As a result, observed consumer inertia can be explained by much smaller costs of switching.

Table 3.4: Parameter estimates in the structural switching cost model with fully attentive consumers

	(1)	(2)	(3)	(4)	(5)	(6)
Intercept (switching cost)	5.245***	5.240***	5.348***	5.532***	5.602***	5.699***
	(0.352)	(0.351)	(0.373)	(0.409)	(0.421)	(0.439)
Campaign	-1.183***	-1.184***	-1.221***	-1.258***	-1.282***	-1.294***
	(0.186)	(0.186)	(0.193)	(0.202)	(0.205)	(0.209)
ν_0	0.293***	0.293***	0.293***	0.293***	0.293***	0.293***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
ν_1	1.416***	1.416***	1.446***	1.494***	1.510***	1.534***
	(0.090)	(0.090)	(0.097)	(0.106)	(0.109)	(0.113)
Insurer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Risk controls		Yes	Yes	Yes	Yes	Yes
Personal controls			Yes	Yes	Yes	Yes
Vehicle controls				Yes	Yes	Yes
Payment controls					Yes	Yes
Contract channel controls						Yes
Observations	6,766	6,766	6,766	6,766	6,766	6,766

10.14754/CEU.2015.04

Table 3.5: Structural estimates are robust to a wide range of price calculation errors

	Error threshold in 2010 contract price calculation						
	±1%	±5%	±10%	±15%	±20%	±30%	±50%
Switching cost	\$52	\$66	\$60	\$64	\$106	\$107	\$93
Attention baseline	29%	29%	28%	29%	34%	35%	33%
Campaign effect	23%	23%	22%	27%	31%	31%	26%
Observations	4,246	6,766	7,827	9,749	13,219	16,140	18,036

reasonable assumptions on unobserved discount eligibility, I can calculate all alternative prices in both periods, but these calculations are not perfectly accurate (see the appendix for more details). Since I know the actual price for the chosen contract in 2010, I can filter the observations by comparing the actual price and the calculated price and dropping the contracts for which the difference is too large.

Table 3.5 shows selected results for various error thresholds in the price calculation for the chosen contract in 2010. The baseline sample corresponds to the 5% column. The rows show the point estimates for the main parameters of interest: the switching cost (without control variables, cf. column (1) of Table 3.2), the baseline attention level, and the effect of the media campaign on the share of attentive people.

The lesson from Table 3.5 is that the estimates are robust to a wide range of price calculation errors. Up to the 15% threshold, practically all numbers are the same (although switching costs tend to be somewhat lower for the 1% sample), whereas the estimates are markedly higher for the 20% threshold and above. Since the baseline specifications relied on the 5% sample, measurement error in prices is not a serious concern regarding the accuracy of the main results.

3.6 Policy counterfactuals

In this section, I evaluate the effect of two hypothetical policy measures on consumer surplus. The first measure is aimed at reducing the cost of searching for alternatives and conducting the contract switch, which in the baseline model is estimated to be around \$65. The second measure is aimed at increasing consumer attention to the switching opportunity.

In practice, the two kinds of interventions cannot be cleanly separated. A switching cost reduction can only have an effect on behavior if it is noticed. But if a policy is designed to be noticed, then it also raises the level of attention. It is, on the other hand, possible to raise attention without affecting the cost of switching: according to the estimates in Table 3.2, for example, the media campaign only increases people's attentiveness, but does not reduce their switching costs.

The policy evaluation only considers the demand side response: do consumers switch more often, and if they do, how much do they gain in the process? The supply side of the market is too complex to model within the scope of this paper, since little is known about the cost structure of providing auto liability insurance, and more importantly, companies are also keen to cross-sell other types of insurance products (with higher margins) to their customers. Without a supply side model that takes into account these considerations, it is hard to provide a quantitative assessment of how insurers would react to more active switching behavior by consumers.

I calculate the change in expected consumer surplus using the same procedure as Hortaçsu et al. (2014), who themselves modify the original method of Small and Rosen (1981) to allow for inattentive decision-makers. Specifically, the expected surplus change for consumer n is:

$$\Delta E\left(CS\right)_n = \log\left(\sum_{j\in J_n^{1,CTF}} e^{V_{nj1}^{CTF}}\right) - \log\left(\sum_{j\in J_n^{1,BL}} e^{V_{nj1}^{BL}}\right) \tag{3.11}$$

where BL stands for the baseline scenario and CTF for the policy counterfactuals. V_{nj1} is the deterministic part of utility provided to consumer n by contract j at $t = 1$ (that is: $V_{nj1} = U_{nj1} - \nu_1 \varepsilon_{nj1}$, where U_{nj1} is given by equation (3.3)).

The choice sets deserve additional explanation, since this is where the surplus calculation with inattentive consumers differs from the original formulation of Small and Rosen (1981). For an attentive person, the choice set contains all alternatives available at the time of the switching decision, whereas for an inattentive person, the choice set contains only one alternative, the contract chosen in the previous period. Whether a person is attentive or inattentive is determined by a random binary draw. With probability θ_n, driver n will be able to pick from all contracts in J^1, and with probability $1 - \theta_n$, he will only be able to "choose" his default contract, simplifying expression (3.11) to:

$$\Delta E\left(CS\right)_n = V_{nj1}^{CTF} - V_{nj1}^{BL} \tag{3.12}$$

Since utility is measured in dollar terms, the interpretation of the values for $\Delta E\left(CS\right)_n$ is also in dollars.

The first hypothetical policy intervention (switching cost reduction) acts by reducing the intercept of βZ_{nj1} in equation (3.3), hence it will make V_{nj1}^{CTF} higher than V_{nj1}^{BL} by the assumed change in switching costs for all the alternatives that involve switching. As for the utility provided by the default alternative, a change in the cost of switching does not matter.

The second hypothetical intervention (increase in attentiveness) does not affect the utility of individual alternatives directly, resulting in $V_{nj1}^{CTF} = V_{nj1}^{BL}$ for all n and j. However, since the policy measure increases the chance of being attentive, it will more often lead to a full choice set under the counterfactual scenario than under the baseline.

The calculation of expected surplus change brings up one additional issue. Since attention is a binary variable that is randomly drawn (with a known success probability), $\Delta E\left(CS\right)_n$ is essentially a simulated outcome that may differ each time it is calculated. I mitigate the sampling error in the procedure by calculating the change in expected surplus for 100 independent draws of the binary attention indicator according to $(\theta_1, ..., \theta_N)$ and averaging the result.

Figures 3.3 and 3.4 show the results of the counterfactual simulations. In Figure 3.3, I vary the cost of switching between \$0 and \$65 on the horizontal axis, and plot the average change in expected consumer surplus on the vertical axis. As the figure shows, the effect is close to linear in the range of costs that I consider. One

10.14754/CEU.2015.04

Figure 3.3: Reducing switching costs by \$65 (to \$0) raises consumer surplus by \$15.6/person

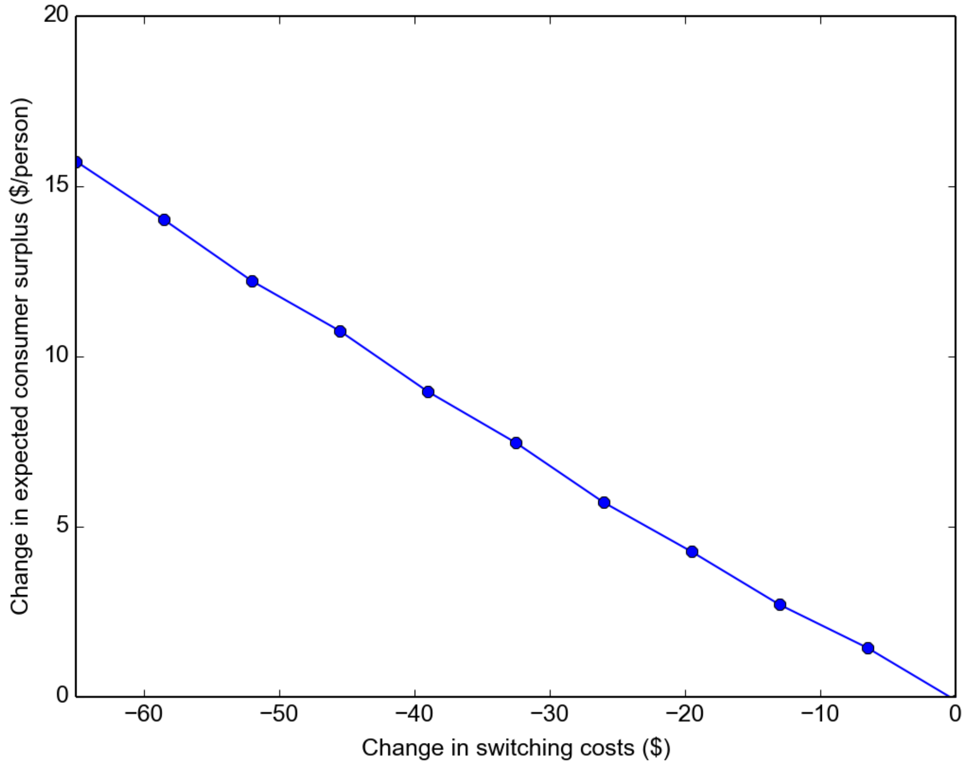

dollar decrease in switching costs increases expected consumer surplus by about 24 cents. The power of financial incentives is mostly lost on those 69% of consumers who pay no attention to the switching opportunity.

In Figure 3.4, the horizontal axis shows the hypothesized increase in attention probability from 0 to 25 percentage points. The maximum of this range is only slightly above the estimated marginal effect of the media campaign on attention probabilities. Again, expected consumer surplus changes linearly with the policy variable. A 10 percentage point increase in attention probability raises consumer surplus by \$7.74/person. Accordingly, the campaign is worth about \$18 to each consumer on average.

The conclusion from the policy counterfactuals is that—at least initially—consumers are much better served by policy measures aimed at increasing their awareness to the switching opportunity than by attempts to decrease the time and effort cost of switching. There is also strong complementarity between the two kinds of interventions: reductions in switching costs are much more effective when a larger share of people consider the contract switching decision.

10.14754/CEU.2015.04

Figure 3.4: A 10 percentage point increase in attention probability raises consumer surplus by \$7.7/person

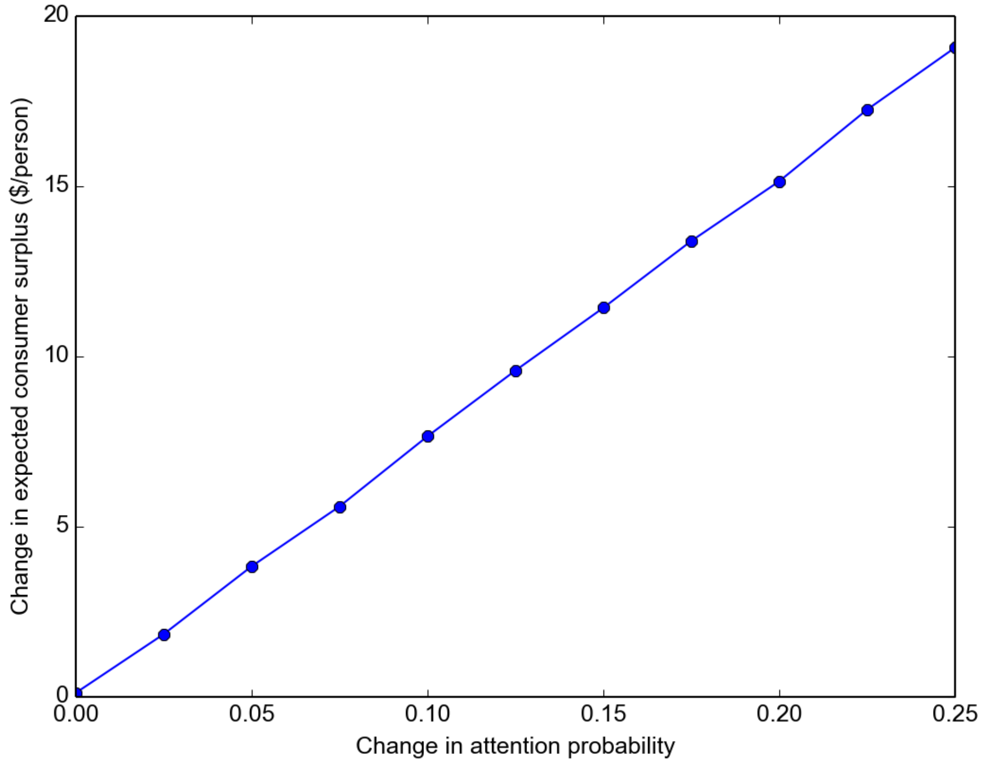

10.14754/CEU.2015.04

3.7 Conclusions

I estimated a structural model on consumer switching, in which switching costs and inattention influence switching decisions through separate channels. I found that inattention to the switching opportunity is widespread: 70 percent of people ignore the decision problem, and the campaign only reaches every third person who would otherwise be inattentive. Estimated mean switching costs are around $65, but the failure to account for inattention would bias these estimates upwards by an order of magnitude.

I have also evaluated two hypothetical policy experiments, one aimed at reducing switching costs, and a second one at drawing people's attention to the switching decision. For plausible changes in the structural parameters, the second intervention proved to be more beneficial to consumers.

The methodology of the chapter to jointly estimate switching costs and inattention could be extended to other choice situations as well. There are a large number of studies that structurally estimate fixed costs associated with stickiness in labor economics, international trade, or monetary macroeconomics. Generally, these studies estimate large fixed costs that prevent people from choosing better alternatives, but much of this inertia could be due to simple inattention to the decision problem. Using the structural approach of this chapter, the bias could be corrected, allowing researchers to draw more accurate conclusions from the evaluation of counterfactual policy measures.

10.14754/CEU.2015.04

Chapter 4

Bibliography

Armstrong, Mark (2006), 'Competition in two-sided markets', *The RAND Journal of Economics* **37**(3), 668–691.

Armstrong, Mark and Jidong Zhou (2011), 'Paying for prominence', *The Economic Journal* **121**(556), 368–395.

Armstrong, Mark, John Vickers and Jidong Zhou (2009), 'Prominence and consumer search', *The RAND Journal of Economics* **40**(2), 209–233.

Artuç, Erhan, Shubham Chaudhuri and John McLaren (2010), 'Trade shocks and labor adjustment: A structural empirical approach', *American Economic Review* **100**(3), 1008–1045.

Baye, Michael R. and John Morgan (2001), 'Information gatekeepers on the internet and the competitiveness of homogeneous product markets', *American Economic Review* **91**(3), 454–474.

Baye, Michael R., John Morgan and Patrick Scholten (2004), 'Price dispersion in the small and in the large: Evidence from an internet price comparison site', *Journal of Industrial Economics* **52**(4), 463–496.

Baye, Michael R., John Morgan and Patrick Scholten (2006), 'Information, search and price dispersion', *Handbook of Economics and Information Systems* pp. 323–375.

Baye, Michael R., Xiaxun Gao and John Morgan (2011), 'On the optimality of clickthrough fees in online markets', *The Economic Journal* **121**(November), 340–367.

Belleflamme, Paul and Eric Toulemonde (2009), 'Negative intra-group externalities in two-sided markets', *International Economic Review* **50**(1), 245–272.

Berger, Lawrence A., Paul R. Kleindorfer and Howard Kunreuther (1989), 'A dynamic model of the transmission of price information in auto insurance markets', *Journal of Risk and Insurance* **56**(1), 17–33.

10.14754/CEU.2015.04

Berndt, E. K., B. H. Hall, R. E. Hall and J. A. Hausman (1974), 'Estimation and inference in nonlinear structural models', *Annals of Economic and Social Measurement* **3**(4), 653–665.

Brown, Jennifer and John Morgan (2009), 'How much is a dollar worth? Tipping versus equilibrium coexistence on competing online auction sites', *Journal of Political Economy* **117**(4), 668–700.

Caillaud, Bernard and Bruno Jullien (2003), 'Chicken & egg: Competition among intermediation service providers', *The RAND Journal of Economics* **34**(2), 309–328.

Chen, Yongmin and Chuan He (2011), 'Paid placement: Advertising and search on the internet', *The Economic Journal* **121**(556), 309–328.

Chetty, Raj, Adam Looney and Kory Kroft (2009), 'Salience and taxation: Theory and evidence', *American Economic Review* **99**(4), 1145–1177.

Cummins, David J., Dan M. McGill, Howard E. Winklevoss and Robert A. Zelten (1974), 'Consumer attitudes toward auto and homeowners insurance', *Department of Insurance, Wharton School, University of Pennsylvania* .

Dahlby, Bev and Douglas S. West (1986), 'Price dispersion in an automobile insurance market', *Journal of Political Economy* **94**(2), 418–438.

Das, Sanghamitra, Mark J. Roberts and James R. Tybout (2007), 'Market entry costs, producer heterogeneity, and export dynamics', *Econometrica* **75**(3), 837–873.

DellaVigna, Stefano and Joshua M. Pollet (2009), 'Investor inattention and Friday earnings announcements', *Journal of Finance* **64**(2), 709–749.

Einav, Liran, Theresa Kuchler, Jonathan Levin and Neel Sundaresan (2014), 'Assessing sale strategies in online markets using matched listings', *American Economic Journals: Microeconomics* **6**(Forthcoming).

Ellison, Glenn and Drew Fudenberg (2003), 'Knife-edge or plateau: When do market models tip?', *The Quarterly Journal of Economics* **118**(4), 1249–1278.

Ellison, Glenn, Drew Fudenberg and Markus Mobius (2004), 'Competing auctions', *Journal of the European Economic Association* **2**(1), 30–66.

Ellison, Glenn and Sara Fisher Ellison (2009), 'Search, obfuscation, and price elasticities on the internet', *Econometrica* **77**(2), 427–452.

Farrell, Joseph and Carl Shapiro (1988), 'Dynamic competition with switching costs', *RAND Journal of Economics* **19**(1), 123–137.

Farrell, Joseph and Paul Klemperer (2007), Coordination and lock-in: Competition with switching costs and network effects, *in* M.Armstrong and R.Porter, eds, 'Handbook of Industrial Organization', Vol. 3, North Holland, chapter 31, pp. 1967–2002.

Finkelstein, Amy (2009), 'E-ZTAX: Tax salience and tax rates', *Quarterly Journal of Economics* **124**(3), 969–1010.

Galeotti, Andrea and José Luis Moraga-González (2009), 'Platform intermediation in a market for differentiated products', *European Economic Review* **53**(4), 417–428.

Golosov, Mikhail and Robert E. Lucas (2007), 'Menu costs and Phillips curves', *Journal of Political Economy* **115**(2), 171–199.

Hałaburda, Hanna and Mikołaj Jan Piskorski (2013), 'Competing by restricting choice: The case of search platforms', *Harvard Business School Working Paper* (10-098).

Handel, Benjamin R. (2013), 'Adverse selection and switching costs in health insurance markets: When nudging hurts', *American Economic Review* **103**(7), 1–48.

Hirshleifer, David, Sonya Seongyeon Lim and Siew Hong Teoh (2009), 'Driven to distraction: Extraneous events and underreaction to earnings news', *Journal of Finance* **64**(5), 2287–2323.

Honka, Elizabeth (2014), 'Quantifying search and switching costs in the U.S. auto insurance industry', *The RAND Journal of Economics* (Forthcoming).

Hortaçsu, Ali, Seyed Ali Madanizadeh and Steven L. Puller (2014), 'Power to choose? An analysis of choice frictions in the residential electricity market', *Unpublished* .

Hossain, Tanjim and John Morgan (2006), '...Plus shipping and handling: Revenue (non)equivalence in field experiments on eBay', *Advances in Economic Analysis and Policy* **5**(2).

Levin, Jonathan D. (2011), 'The economics of internet markets', *NBER Working Paper* pp. 1–33.

Luco, Fernando (2014), 'Distinguishing sources of inertia in a defined-contribution pension system', *Unpublished* .

Miravete, Eugenio J. and Ignacio Palacios-Huerta (2014), 'Consumer inertia, choice dependence, and learning from experience in a repeated decision problem', *Review of Economics and Statistics* **96**(3), 524–537.

Moraga-González, José Luis and Matthijs R. Wildenbeest (2012), 'Comparison sites', *The Oxford Handbook of the Digital Economy* pp. 224–253.

Rochet, Jean-Charles and Jean Tirole (2003), 'Platform competition in two-sided markets', *Journal of the European Economic Association* **1**(4), 990–1029.

Scale Research (2010), 'Consumer expectations in the mandatory auto liability insurance market in Hungary with non-synchronized switching', *Study prepared at the request of the Hungarian Competition Authority* (in Hungarian).

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

10.14754/CEU.2015.04

Train, Kenneth E. (2003), *Discrete Choice Methods with Simulation*, Cambridge University Press.

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

Viard, V. Brian (2007), 'Do switching costs make markets more or less competitive? The case of 800-number portability', *The RAND Journal of Economics* **38**(1), 146–163.

Zhang, Tianle (2009), 'Price-directed consumer search', *Working Paper, University of Colorado* (09-08).

10.14754/CEU.2015.04

Appendix A

Appendix for Chapter 1

10.14754/CEU.2015.04

1.1 Proof of Proposition 1.1

The following series of lemmas show that no pure strategy equilibrium exists among advertising firms on a given price comparison site. Among mixed strategy equilibria, I will concentrate on those involving symmetric strategies by sellers, for which I derive a number of features.

Lemma A.1. *If a platform charges positive listing fees* $(a > 0)$*, then advertising firms never charge prices at, or below, the effective marginal cost of the product* $m + \frac{c}{\beta}$*, or above the reservation price* (r) *of consumers.*

Proof. Setting a price below $m + \frac{c}{\beta}$ yields negative profits, whereas charging $p = m + \frac{c}{\beta}$ or $p > r$ yields zero (gross) profits to the firm, in addition to which a positive listing fee must be paid to the platform. Since advertising firms cannot do worse than the outside option of zero (net) profits, $p \leq m + \frac{c}{\beta}$ and $p > r$ cannot be part of an equilibrium pricing strategy. □

Lemma A.2. *There is no equilibrium in pure pricing strategies within the support of* $\left(m + \frac{c}{\beta}, r\right)$ *on a given platform.*

Proof. Suppose that seller i charges a price p_i and is placed kth on the list of offers. If seller i is not tied with any other seller, then it can increase its expected profits by charging a slightly higher price, since its expected number of transactions do not change as long as the price increase does not result in a tie or in a place swap. On the other hand, if seller i is tied with $l-1$ other sellers, then—assuming that there are $\hat{\mu}$ "incoming" consumers to the tied firms—seller i gets $\beta \frac{1+(1-\beta)+(1-\beta)^2+...+(1-\beta)^{l-1}}{l} \cdot \hat{\mu}$ buyers on average, which is strictly less than it could get by reducing its price very slightly $(\beta\hat{\mu})$. As a result, it is never a best response to stay in a tie with others, which has a similar flavor to the standard argument of Bertrand competition. □

Lemma A.3. *There is no symmetric mixed strategy equilibrium in which firms place a positive probability mass on any price in the interval* $\left(m + \frac{c}{\beta}, r\right]$.

Proof. Suppose, on the contrary, that there exists a price $p \in \left(m + \frac{c}{\beta}, r\right]$ that is charged with a probability $q > 0$ in a symmetric mixed strategy equilibrium. Then, with a positive probability at least two sellers would tie at this price, resulting in a sharing of potential buyers (and expected profits from these buyers). If one firm very slightly reduced the price that it charges with probability q from p to $p - \varepsilon$, it would receive a discontinuous positive jump in expected profits (the amount which was formerly shared with others in case of a tie), but only lose an infinitesimally small (order ε) amount of expected profit in all other cases. Therefore, charging p with a strictly positive probability q cannot be a best reply to the same strategy for any $p \in \left(m + \frac{c}{\beta}, r\right]$. □

10.14754/CEU.2015.04

We have thus shown that with positive listing fees, only mixed strategy advertising and pricing equilibria exist, and if we are willing to restrict ourselves to symmetric ones among these, then the distribution from which prices are drawn must be continuous and cannot include probability "masses" on any price within the interval $\left(m + \frac{c}{\beta}, r\right]$.

In the next claim, I point to a further important price setting feature.

Lemma A.4. *The reservation price of consumers is always included in the support of the symmetric equilibrium pricing strategy.*

Proof. Suppose that the highest price charged in equilibrium is $\overline{p}$, that is, $\Pr\left(p < \overline{p}\right) = 1$ and $\Pr\left(p < \hat{p}\right) < 1$ for any other $\hat{p} < \overline{p}$ that is part of the equilibrium strategy. When charging $\overline{p}$, a firm can only make sales to people who are considering the firm as a last possible choice on the platform. In this situation, the number of consumers served is independent of the price chosen, as long as the price does not exceed the reservation value r. Therefore, $\overline{p}$ should be set to maximize profits with an exogenous number of sales, resulting in $\overline{p} = r$. □

1.2 Proof of symmetric equilibrium existence and uniqueness without per-click charges

Let us characterize the symmetric equilibria of the game. That is, set $n_1 = n_2 \equiv n$, which also results in equal market shares, and solve the first order condition of platform profit maximization in (1.32):

$$h\left(n\right) \equiv \left\{\left|\log\left(1-\beta\right)\right| + \left|\log\left(1-\beta\right)\right| \frac{\beta}{1-\beta} \cdot n - \frac{\beta}{1-\beta}\right\} \left(1-\beta\right)^{n} \left(\frac{r-m}{2\sigma}\right) + \frac{1}{n} - \left|\log\left(1-\beta\right)\right| = 0$$

The function $h\left(n\right)$ is several times continuously differentiable on the positive real line. Its limits are:

$$\lim_{n\to 0^{+}} h\left(n\right) = +\infty$$

$$\lim_{n\to\infty} h\left(n\right) = -\left|\log\left(1-\beta\right)\right| < 0$$

Therefore, $h\left(n\right)$ has at least one root. One symmetric equilibrium surely exists.

If $h\left(n\right)$ were strictly decreasing, the symmetric equilibrium would be unique. However, for some parameter values, there are intervals on the domain of h where the function is increasing. This raises the possibility of h having multiple roots, and therefore the existence of multiple equilibria.

Where might these roots lie? Since the first component of h is positive (by the positivity of marginal surplus when $n \geq 1$), all roots must satisfy:

$$n > \frac{1}{\left|\log\left(1-\beta\right)\right|}$$

which causes the sum of the second and third component of h to be negative.

Moreover, we know that only one root can exist on the domain where h is strictly decreasing:

$$h'(n) = |\log(1-\beta)| \left\{ \frac{2\beta}{1-\beta} - |\log(1-\beta)| - |\log(1-\beta)| \frac{\beta}{1-\beta} \cdot n \right\} \cdot (1-\beta)^n \left(\frac{r-m}{2\sigma} \right) - \frac{1}{n^2} < 0$$

The inequality holds if

$$\frac{2\beta}{1-\beta} - |\log(1-\beta)| - |\log(1-\beta)| \frac{\beta}{1-\beta} \cdot n \leq 0.$$

Rearrangement yields:

$$n \geq \frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta}$$

Of course, h might also be strictly decreasing for some $n < \frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta}$, but our concern here is to establish the interval on which multiple roots must exist if there are multiple equilibria at all.

One caveat at this point: from an economic point of view, equilibrium outcomes are only interesting if the number of sellers in equilibrium is at least 1. This rules out all situations for which $\frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta} < 1$. That is, for the proof of uniqueness regarding economically interesting symmetric equilibria, we only have to be concerned with cases in which $\beta \lessapprox 0.7968$ holds ("admissible" values of β).

Summarizing what we have found so far: (1) h crosses the horizontal axis from above at least once; (2) all crossings are at $n > \frac{1}{|\log(1-\beta)|}$; (3) if there are multiple equilibria, h must either "touch" the horizontal axis from above at least once (i.e. have a local minimum that equals 0), or cross the horizontal axis from below at least once; (4) for all $n \geq \frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta}$, h is strictly decreasing, and therefore can only cross the horizontal axis from above (if at all); and (5) we only need to be concerned with parameter values in which $\beta \lessapprox 0.7968$.

Consequently, roots with the "touching" or "crossing-from-below" property must lie in the non-empty open interval:

$$I \equiv \left(\frac{1}{|\log(1-\beta)|}, \frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta} \right).$$

Moreover, if a root with a "crossing-from-below" property exists, then another one crossing from above must also exist within the same interval, to the left of the former. Therefore, (by continuous differentiability) the following joint condition is necessary for the existence of multiple symmetric equilibria: there has to be at least one point n in the open interval I for which $h'(n) = 0$ and $h(n) \leq 0$ hold at the same time. That is, a non-positive local minimum must exist on I.

Expanding the two conditions:

$$h'(n) = |\log(1-\beta)| \left\{ \frac{2\beta}{1-\beta} - |\log(1-\beta)| - |\log(1-\beta)| \frac{\beta}{1-\beta} \cdot n \right\} \cdot (1-\beta)^n \left(\frac{r-m}{2\sigma} \right) - \frac{1}{n^2} = 0 \quad \text{(A.1)}$$

and

$$h(n) = \left\{ |\log(1-\beta)| + |\log(1-\beta)| \frac{\beta}{1-\beta} \cdot n - \frac{\beta}{1-\beta} \right\} \cdot (1-\beta)^n \left(\frac{r-m}{2\sigma} \right) + \frac{1}{n} - |\log(1-\beta)| \leq 0 \quad \text{(A.2)}$$

Expressing $(1-\beta)^n \left(\frac{r-m}{2\sigma}\right)$ from (A.1) and substituting it into (A.2):

$$h(n) = \frac{1 + \frac{\beta}{1-\beta} \cdot n - \frac{1}{|\log(1-\beta)|} \frac{\beta}{1-\beta}}{\left[\frac{2\beta}{1-\beta} - |\log(1-\beta)| - |\log(1-\beta)| \frac{\beta}{1-\beta} \cdot n \right] n^2} + \frac{1}{n} - |\log(1-\beta)| \leq 0$$

The denominator of the first fraction is positive, whenever $n < \frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta}$. The numerator is greater than 1 when $n > \frac{1}{|\log(1-\beta)|}$, and increasing in n. Both of these conditions are true on I.

Moreover, the denominator is decreasing in n when:

$$\frac{\partial}{\partial n} \left[\frac{2\beta}{1-\beta} n^2 - |\log(1-\beta)| n^2 - |\log(1-\beta)| \frac{\beta}{1-\beta} \cdot n^3 \right] < 0$$

which is true for:

$$n > \frac{2}{3} \left[\frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta} \right].$$

Within I, $n > \frac{1}{|\log(1-\beta)|}$, which satisfies the monotonicity condition when:

$$\frac{2}{\beta} - 2 \geq \frac{1}{|\log(1-\beta)|}$$

This inequality holds for all admissible values of $0 < \beta \leq 1 - \frac{1}{e}$. Therefore, the large fraction in $h(n)$ is always larger than:

$$\frac{|\log(1-\beta)|^2}{\frac{\beta}{1-\beta} - |\log(1-\beta)|}$$

whereas

$$\frac{1}{n} - |\log(1-\beta)| > \frac{1}{\frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta}} - |\log(1-\beta)|$$

on I. Thus:

$$h(n) > \frac{|\log(1-\beta)|^2}{\frac{\beta}{1-\beta} - |\log(1-\beta)|} + \frac{1}{\frac{2}{|\log(1-\beta)|} - \frac{1-\beta}{\beta}} - |\log(1-\beta)| \quad \text{(A.3)}$$

10.14754/CEU.2015.04

The left hand side of (A.3) is positive on the interval $(0, 0.877)$ which contains all admissible values of β.

In summary, whenever $h'(n) = 0$, there is no n within I, for which $h(n) \leq 0$ would hold true. Consequently, the platforms' profit maximizing game has a unique symmetric equilibrium.

1.3 Numerical evidence for the non-existence of asymmetric equilibria without per-click charges

In the previous proof, I have shown that a unique symmetric equilibrium always exists when platforms are unable to charge per-click fees. What remains to be proven is that there are no asymmetric equilibria, which might be interpreted as market tipping.

The implicit best response functions in (1.32) for $n_1 \neq n_2$ are algebraically complex and do not yield to analytic proofs, unlike in the symmetric case shown above. Instead, I perform a large number of best response simulations over a range of parameters to search for signs of multiple equilibria. Since there is only one symmetric equilibrium, the existence of multiple equilibria necessarily means that some of these involve asymmetric outcomes (i.e. one platform having more buyers and sellers than the other). The simulation method is the following.

First, I choose a range of parameter values and the distribution of sampling points for β, σ, and $r - m$.[1] Second, I choose a range and distribution for sampling seller numbers.[2] Third, for all possible combinations of the parameters, I establish that equation (1.32) has a unique solution for any n_i in my sample, provided that n_j is also restricted to be above 1. Finally, I calculate the following expression for all parameter combinations:

$$d(n_i) \equiv BR_i [BR_j (n_i)] - n_i \tag{A.4}$$

where BR_j denotes the best response of platform j to the number of sellers listed by platform i, calculated by solving equation (1.32) numerically (and enforcing $BR_j (n_i) \geq 1$).

When $d(n_i) = 0$, n_i is part of an equilibrium play. Therefore, showing the non-existence of multiple equilibria (and hence of asymmetric equilibria) numerically is equivalent to showing that $d(n_i)$ crosses the horizontal axis at most once.

Using the sampled points, I have confirmed that $d(n_i)$ indeed only has one root in all cases. Figures A.1-A.3 show the shape of $d(n_i)$ graphically for various cross-sectional views of the parameters. The function is mostly downward sloping, but sometimes exhibits interesting turns (hence the difficulty in the analytic approach).

To conclude, I have shown numerically that for a wide range of parameter values, it is very likely that the platform competition model without per-click charges has no asymmetric equilibria, and hence the unique symmetric equilibrium found above is also a unique equilibrium overall.

[1] $0 < \beta < 1$, $0.1 \leq \sigma \leq 100$, and $1 \leq r - m \leq 1000$. I have 20 sampling points for β, 50 for σ, and 40 for $r - m$. Sampling points are distributed uniformly: on a linear scale for β and on log scales for σ and $r - m$.

[2] Sellers are sampled at 100 points, uniformly distributed on a log scale between 1 and 100.

Figure A.1: Best response uniqueness with various parameter values

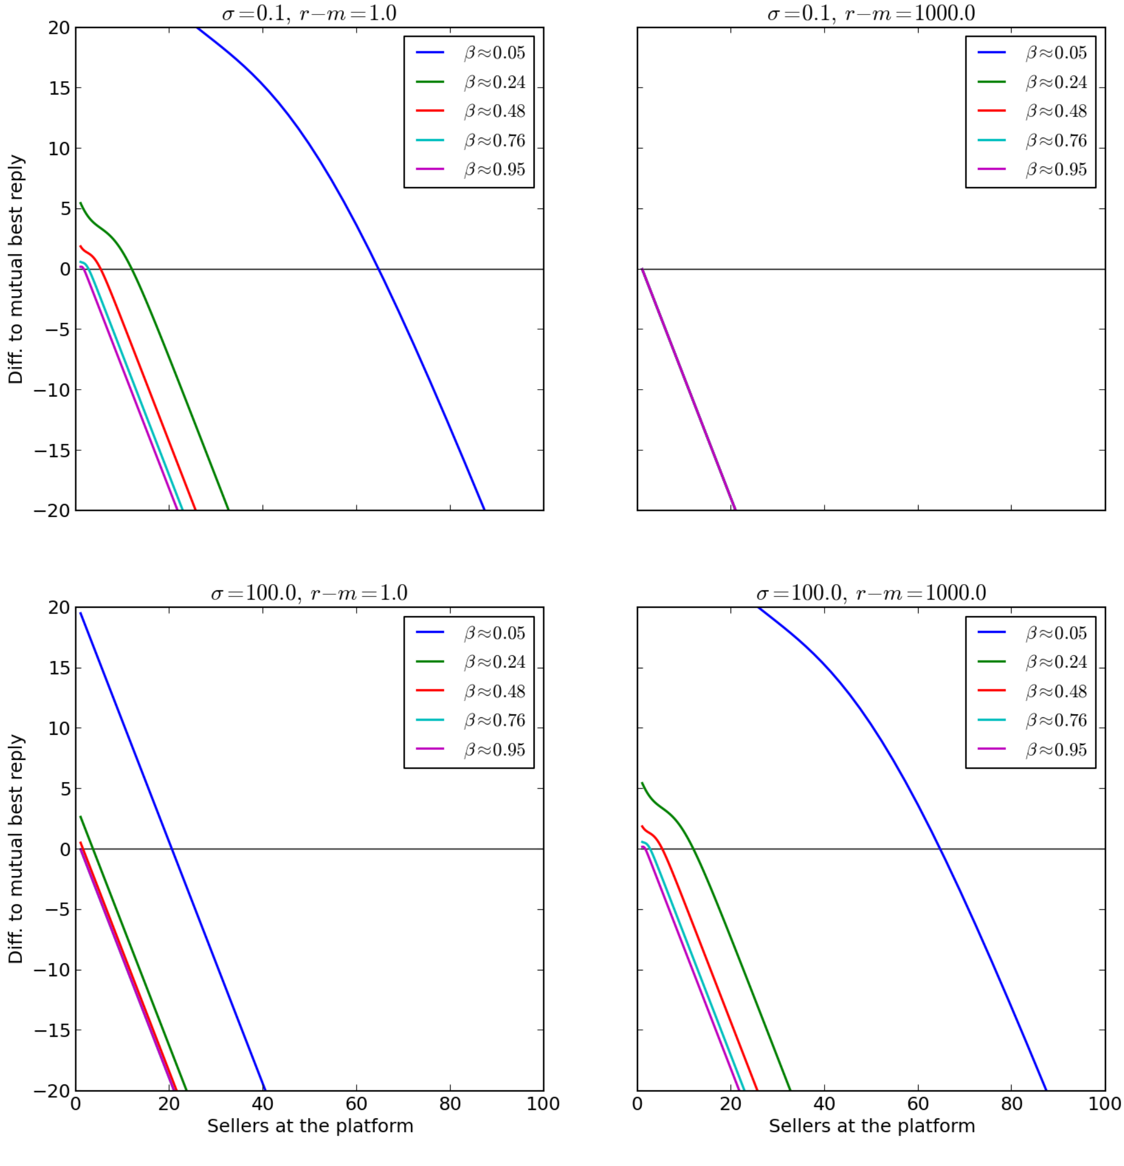

Figure A.2: Best response uniqueness with various parameter values

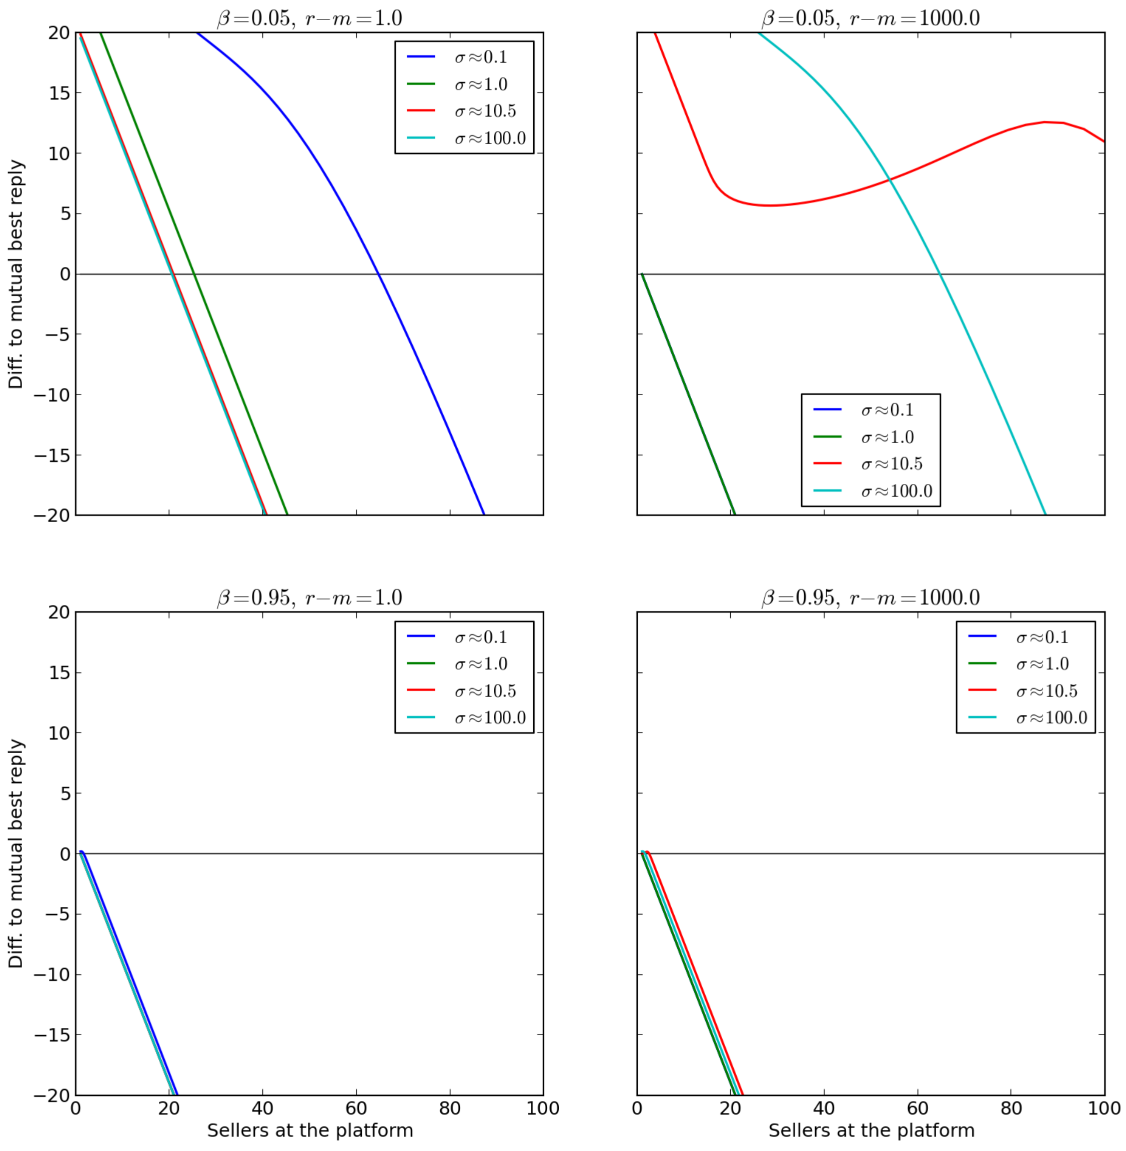

Figure A.3: Best response uniqueness with various parameter values

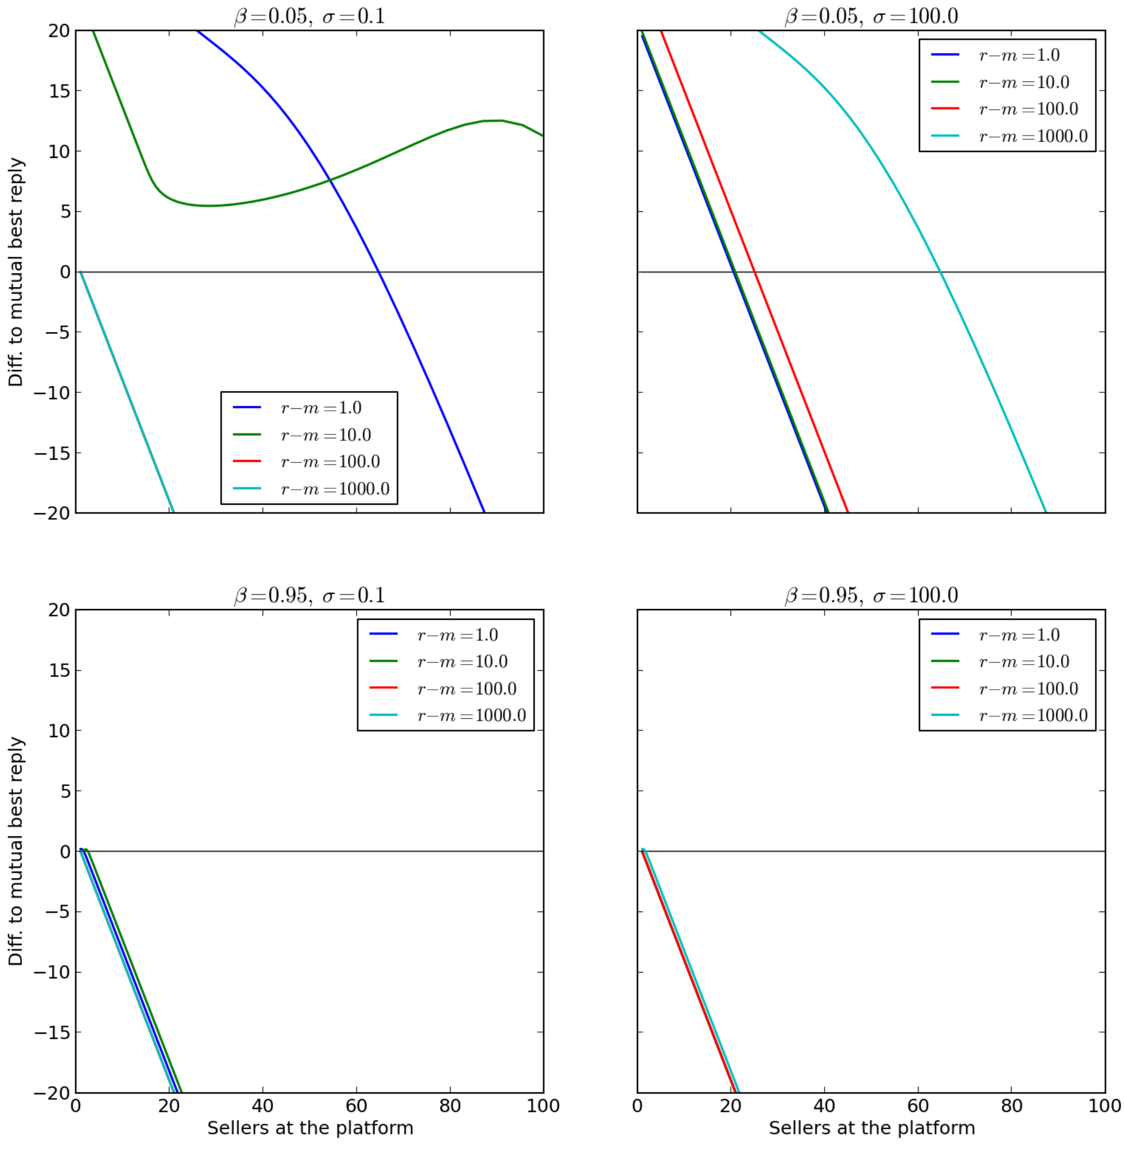

10.14754/CEU.2015.04

1.4 Proof of the existence of asymmetric equilibria with per-click charges and fixed cost of price advertising

When sellers incur a positive fixed cost κ each time they list prices on a platform, the first order condition (1.20) of profit-maximization for platforms changes to:

$$\frac{\partial \Pi_j}{\partial n_j} = -\log(1-\beta)\,\mu_j\,(1-\beta)^{n_j}\,(r-m) - \kappa = 0 \tag{A.5}$$

wheras (1.21) remains the same:

$$\frac{\partial \Pi_j}{\partial u_j} = \frac{\partial \mu_j}{\partial u_j}\left\{\left[1-(1-\beta)^{n_j}\right](r-m) - u_j\right\} - \mu_j = 0$$

Writing out the first order conditions for both platforms, and substituting the consumer behavior equations (1.2) and (1.3) yields, after algebraic manipulations:

$$\left[\frac{\kappa}{|\log(1-\beta)|} - \sigma\right]\left[\exp\left(\frac{\Delta u}{\sigma}\right) - \exp\left(-\frac{\Delta u}{\sigma}\right)\right] = \Delta u \tag{A.6}$$

where $\Delta u = u_2 - u_1$, the difference between consumer utility provided by the two platforms. It is immediate to see that a symmetric outcome ($\Delta u = 0$) always forms an equilibrium. The question is, when do asymmetric equilibria exist?

In an asymmetric equilibrium, $\Delta u \neq 0$. First, observe that if $\Delta u = \overline{u}$ satisfies (A.6), then $\Delta u = -\overline{u}$ does as well. In effect, the asymmetric equilibria exist "symmetrically", if at all, which is as we would expect it. From now on, assume that $\Delta u > 0$.

Graphically, the left hand side of (A.6) resembles that of a plain cubic function of Δu with no quadratic, linear, or constant terms. Since the right hand side of (A.6) is positive, a necessary condition for the existence of asymmetric equilibria is:

$$\frac{\kappa}{|\log(1-\beta)|} - \sigma > 0. \tag{A.7}$$

The slope of the left hand side is given by:

$$\left[\frac{\kappa}{\sigma\,|\log(1-\beta)|} - 1\right]\left[\exp\left(\frac{\Delta u}{\sigma}\right) + \exp\left(-\frac{\Delta u}{\sigma}\right)\right]$$

Since this slope is increasing with Δu, the existence of asymmetric equilibria also requires that the slope is less than 1 at $\Delta u = 0$:

$$\frac{\kappa}{\sigma\,|\log(1-\beta)|} < \frac{3}{2}. \tag{A.8}$$

Putting (A.7) and (A.8) together, we get:

$$\sigma\,|\log(1-\beta)| < \kappa < \frac{3}{2}\sigma\,|\log(1-\beta)| \tag{A.9}$$

In summary, platform competition with fixed advertising costs on the sellers' part always has a symmetric equilibrium, and also has two asymmetric equilibria when

10.14754/CEU.2015.04

Figure A.4: Existence of asymmetric equilibria with various parameter values

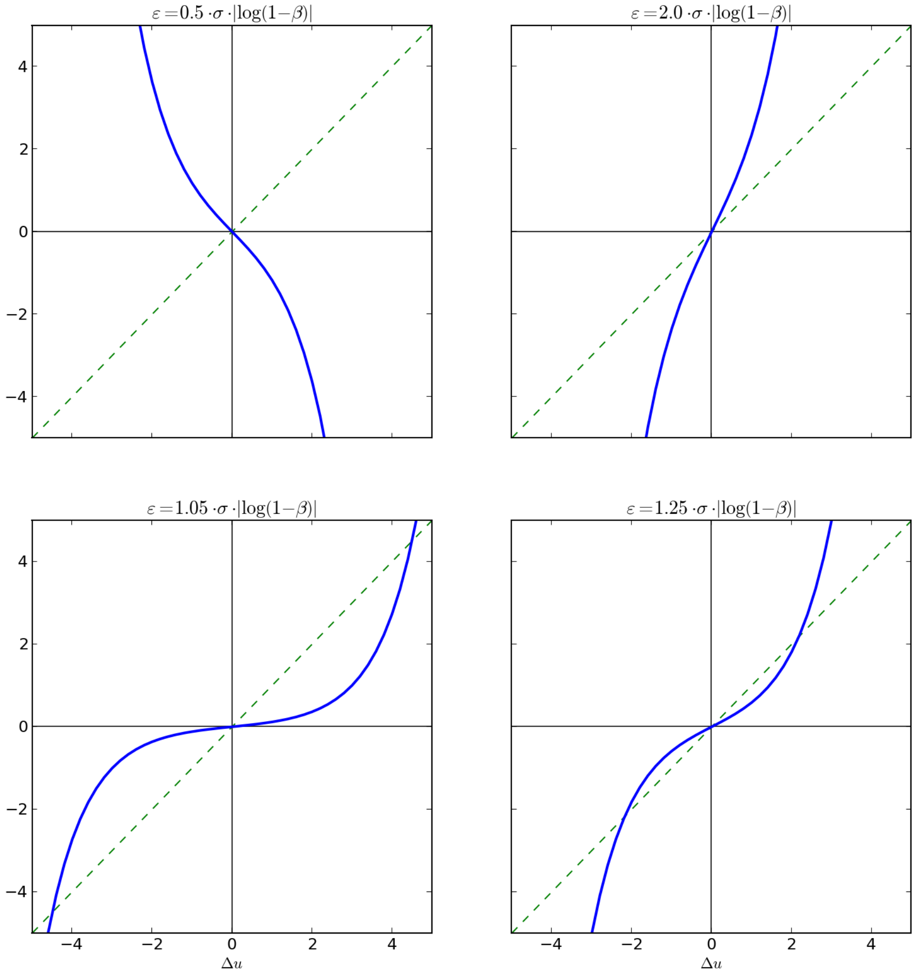

10.14754/CEU.2015.04

the fixed costs are in the range specified by (A.9). In particular, for given values of σ and β, the equilibrium is more asymmetric when κ is closer to its lower boundary. Figure A.4 provides a graphical illustration of the possible outcomes.

When platform j provides more utility to consumers than platform i in equilibrium, it also has a larger market share on the buyers side and lists more sellers than its rival.

10.14754/CEU.2015.04

Appendix B

Appendix for Chapter 2

10.14754/CEU.2015.04

2.1 Data description

I received access to the contract level database of a mid-sized insurance broker in Hungary. In this section, I describe the cleaning and transformation procedures I performed to arrive at my estimation sample.

Originally, the database contained 363,404 contracts, started between 2007 and 2013. For reasons detailed in the main text, my identification procedure only uses contracts started in 2010, which initially numbered 60,286.

The data contained a considerable amount of miscoding (e.g. of birthdates, vehicle characteristics, etc.) Wherever these could be unambiguously corrected, I did so. For the rest of the data, I dropped the record entirely when the miscoded variable was important, otherwise just registered a missing value.

I further restricted my sample to contain personal vehicles (i.e. cars) only, insured for personal use (and not as taxi or a rental vehicle, for example). I dropped all drivers who had more than one vehicle insured to their name. Arguably, these people had different incentives for choosing insurers and tracking contracts than everyday drivers.[1]

For partly similar reasons, I dropped contracts with the insurance company AIM. AIM was a successor of a troubled insurer (TIR) that went out of business by 2010. However, one year later AIM ceased its operations as well, forcing its customers to switch insurers whether they wanted to or not. Since the market share of AIM in the sample is tiny (well below 1 percent), its exclusion is not noticeable on the results.

I also excluded contracts in the M1–M4 risk categories. These drivers have been found at fault in recent accidents. Since I have no information on the date of the damage claims, the estimation of the alternative insurance premia would contain considerable inaccuracy, and hence worsen the quality of my estimates.

One of the main discriminating factors in contract pricing is the power of the vehicle's engine, therefore I took special care to exclude potentially miscoded values (beyond simply checking for numerical correctness). The main source of confusion was that both the cylinder volume and the power were recorded, measured in ccm and kW, respectively. In a non-trivial number of cases, people mixed up the two, or only recorded one, but in the wrong place. Some values were also suspiciously low or high. In the end, I opted for excluding all cylinder volumes below 500 ccm and above 5,000 ccm.

In addition, I checked that vehicle power (in general mainly determined by cylinder volume) is not an outlier given the size of the engine. Figure B.1 plots the volume-power relationship seen in the data. I chose to exclude all records that fell outside the red triangle.[2]

A further sample restriction criterion was based on the rules of contracting. According to the regulation, drivers could only switch contracts within the calendar year (i.e. not on January 1st) if the contract started on 1/2/2010 or later, and even then they have to wait for 365 days for the switching opportunity (at least until 1/2/2011). In other words, contracts starting between 1/2/2010 and 12/31/2010

[1] The number of simultaneously insured vehicles by the same person (sometimes over half a dozen) points towards a gaming of the system.

[2] Note that the scale of the vertical axis hides most of the data points for which the volume is incorrectly recorded twice (for the second time as power).

Figure B.1: Valid and excluded vehicle engine characteristics

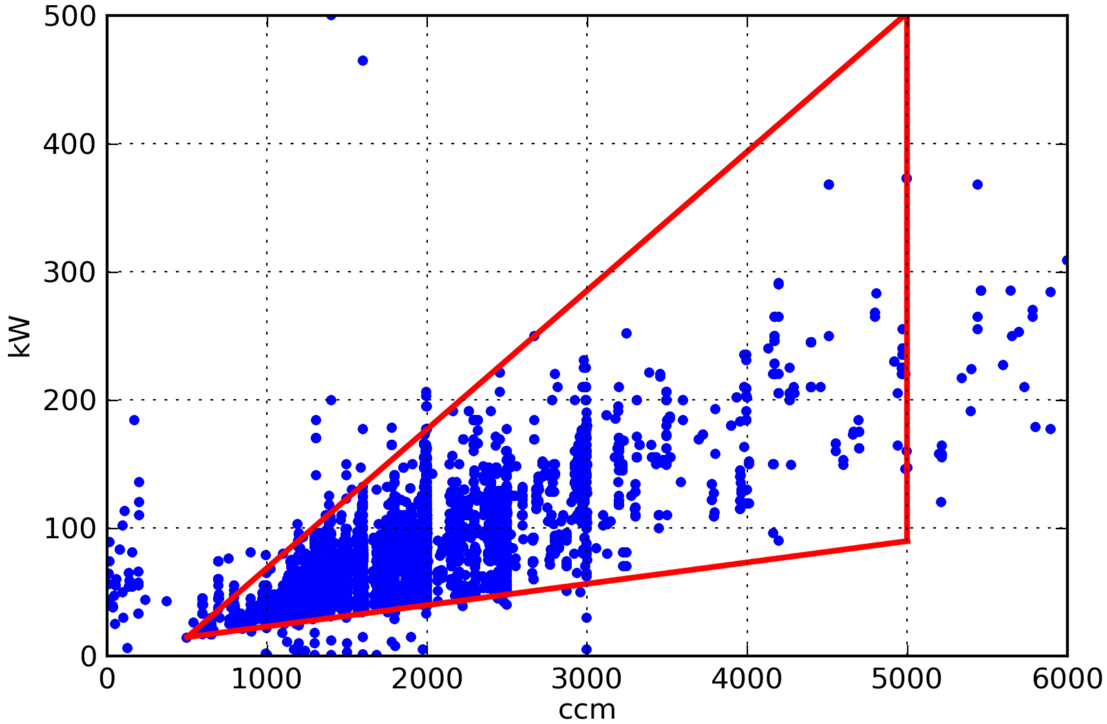

must have a contracting reason *other than* normal contract switching.

Nevertheless, a few dozen contracts during January–March 2010 (but not on 1/1/2010) did denote regular switching as the reason for contracting. Presumably, these belonged to drivers who terminated their existing contracts in time (during November 2009), but for some reason did not sign a new one until 1/1/2010. Hence they were liable to pay a non-coverage penalty on a daily basis, which far exceeded the normal insurance premia. Once they noticed their mistake, they signed the new contract. I still excluded them from my estimation, however, because these people were drivers from the "old regime" with potentially different observed and unobserved characteristics than the January and the February-November group, to which they now belonged in terms of switching periods. Again, the exclusion did not noticeably affect the estimates.

Finally, I dropped all contracts that had a deletion date of 9/24/2012 in the insurance broker's database. An unusually large number of contracts have been deleted on this—otherwise not particularly important—day. There is a practical explanation: the broker company cleaned its database of all contracts that became inactive at some earlier time without being registered as such. Since the contract deletion date is the only information that lets me observe switching and these contracts have uncertain deletion dates, I excluded them.

Overall, the data cleaning resulted in a full sample of 43,692 contracts, out of the original 60,286 starting in 2010.

As described in the robustness section of the Appendix, I used various parts of

10.14754/CEU.2015.04

the full sample for my estimations, based on the accuracy with which I could predict the insurance premium of a given contract in 2010.

Accordingly, I had three different samples: one with ±1%, one with ±5%, and one with ±50% prediction errors. The sample sizes were 13,670, 17,876, and 42,801, respectively, including the contracts that were discontinued before the date of the first anniversary. Excluding the first year dropouts, the contracts that started for a reason other than vehicle acquisition, as well as the *Calendar* and *December* contracts, I had 4,246, 6,766, and 18,036 observations in each sample, the second of which was used for all of my baseline estimations. Naturally, the larger samples nest the smaller ones.

2.2 Establishing contract switching

In the insurance broker's database, contracts refering to the same person and vehicle are not linked over time. However, drivers have unique identifiers and I see detailed information about vehicles, which would enable me—in theory—to track the contracts of a person over time and hence identify when and to which insurer switching occured.

Although this method is somewhat imperfect in itself (similar cars might be mistaken to be the same car), where it really breaks down is the high dropout rate from the sample. Since it is relatively costless to switch to a new broker, or even to bypass them if one has already chosen an insurer, people often switch contracts with the help of a different broker's online interface. In this case, the only feedback the initial broker receives is a message from the driver's insurance company that the contract has been terminated. The termination date (as reported by the insurer) is recorded, but the reason for termination is not.

In theory, contract termination might occur on any day. Contracts are closed down when a vehicle is sold or taken out of traffic, or payments are late by more than 60 days, for example. On the other hand, termination by regular contract switching can only be dated exactly one year (or two, three, etc. years) after the starting date of the contract. Since I observe both the start and the end dates, I can conclude with relative certainty whether the contract was closed down because of switching, or for some other reason.

Specifically, I examine the distribution of the recorded contract lengths, defined as the termination minus the starting date, and look for bunching around 365 days.[3] Figure B.2 illustrates the method for contracts started in 2010.

As the figure shows, the contract length distribution is relatively uniform, save for the region around 365 days.[4] For calendar contracts, the exceptional days are 365 and 370, whereas for intrayear contracts, days 363–366 stand out. Accordingly, I denoted these contracts as being terminated because of switching. If a contract is shorter than 365 (or 363) days, I consider it a "before-switching dropout". If it is longer, the person is a "non-switcher" at the first contract anniversary.

[3] I allow for some leeway in the recording of the termination date.

[4] In fact, there is also some bunching at whole months within the year (hidden by the scale of the horizontal axis), which is most likely an artifact of reporting frequencies.

Figure B.2: Contract length distributions around 365 days for contracts started in 2010

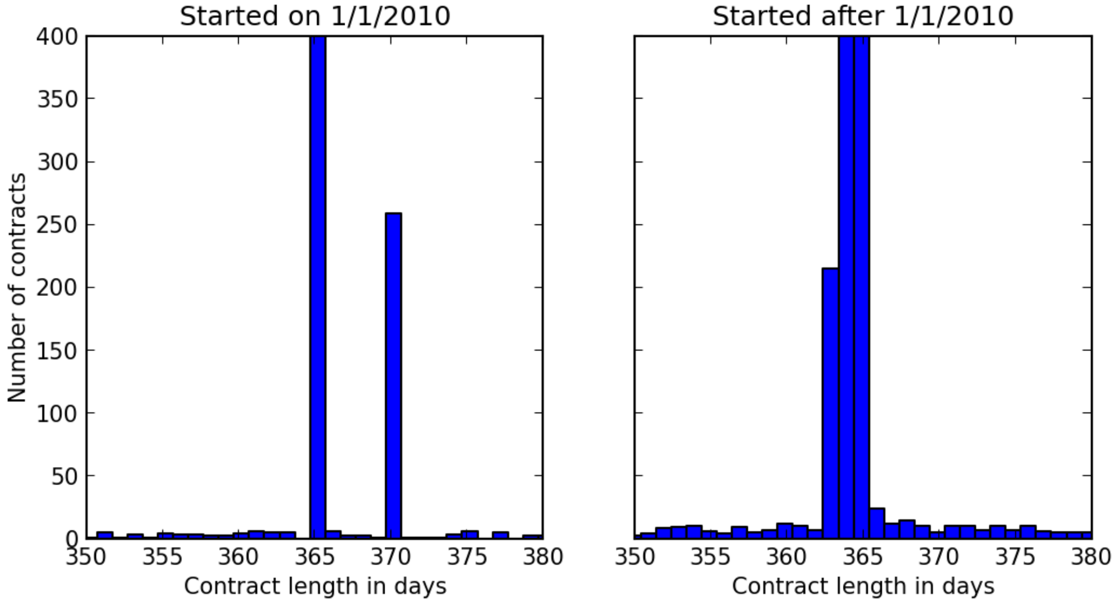

I realize that this method is imperfect. It is quite possible that some switched contracts had a length of 368 days, for example, or that some cars were sold on the 365th day. There are two reasons, however, why the error is unlikely to matter. First, the number of mis-classified contracts is relatively small: if the scale of the vertical axis on Figure B.2 were enlarged to show all switchers, the non-switchers would simply become invisible. Second, the mis-classification is likely to affect January and non-January contracts the same way, and hence cancel out in the comparison of the two groups.

2.3 Insurance fee calculation

My dataset only includes the insurance fee in the first year of an insurance contract. I do not know about the pricing of alternative offers that have not been chosen, and neither do I have information on the actual fees in subsequent years. Therefore I do not observe the monetary benefit (loss) of a (non-)switching decision directly.

To estimate switching costs from observed choices, I need a good predictor of switching benefits. I use the publicly available pricing tables of insurance companies, together with relevant driver- and car-specific information, to make a best guess of what the unobserved prices might be. In this part of the data appendix, I explain my estimation algorithm in detail.

As I have mentioned in the main text, many price setting algorithms are complex, partly relying on personal information that is missing from my database. These missing piece of data usually refer to eligibility of special discounts, many of which are company-specific, and some of which are common across several (if not all) companies.

Table B.1: Input data for insurance pricing algorithms

Status	Usage	Examples
Observed	Common	Driver age, residence, risk category; vehicle power; contract start date, payment method and frequency
	Rare	Year of driver's license, retirement status; vehicle age, fuel, brand
Unobserved	Common	Fault history, driver has young children
	Fairly common	Driver is public servant, communication means (email, mobile)
	Company-specific	Other insurance-contract with same company, special coupons, employment

Table B.1 provides an overview of the type and availability of data used for the determination of the insurance premia. As a rule, discounts might run up to 20-30 percent of the baseline price, but often in a non-monotonic manner. For example, there might be several discount categories worth 5-15 percent of the baseline price, but the overall discount is capped at 30 percent. Similarly, penalties for recent faulty driving might increase the baseline price by 30-50 percent.

This pricing structure makes it impossible to "reverse-engineer" the applied discounts even from observed fees in most cases. For example, a company might offer a 10 percent discount to both those with young children, and to those who are public servants. Observing a 10 percent discount with such a company does not allow me to infer whether a different company that only recognizes young children, but not public servants, would offer a discount to the driver. Moreover, caps might further mask the type of applied discounts.

2.3.1 Calculation procedure

I estimated insurance fees for all drivers and all insurance companies in two years: 2010 and 2011. 2010 is the year when all contracts were signed initially, whereas potential switching decisions were first made in 2011. My focus is on explaining switching behavior in 2011, for which I only need the fees in that year. However, comparing my estimates with actual choices and prices in 2010 help me reduce the effect of the measurement error I introduce by imperfectly proxying prices.

Technically, the estimation involved the programming of insurance price calculators similar to those used by online insurance brokers.[5] To arrive at price estimates, I used all available information in my database, and made "educated guesses" to what I did not observe.

Specifically, I assumed that drivers took advantage of discounts applied to communication means, which amounted to 2-10 percent of the baseline price. Essentially, this means that they were willing to provide their email addresses and mobile phone numbers to the insurance companies. I also assumed that no penalties were applied

[5]Considering that the calculations involved historic insurance fees, I had no opportunity to use currently operating price calculators.

10.14754/CEU.2015.04

based on recent driving history (usually the past 3-5 years).

In addition, I calculated separate prices assuming the presence of young children as well as public servant status. Although my default case included neither, I used the alternative assumptions for robustness checks on my estimates.

All other—mostly company-specific—discounts were set to zero during the price calculation procedure.

2.3.2 Validation of calculated fees

Since have information on insurance premia at the chosen insurer in 2010, I can compare the result of my price calculation to actual data to get a partial idea about the measurement error inherent in the process.

For each contract in the sample, I calculated the proportional estimation error as the difference between the estimated and the actual insurance fee divided by the actual insurance fee. Table B.2 shows the distribution of these errors for various subsamples.

Each row in the table corresponds to the values of the given distribution percentiles in the different samples. A positive number means that I estimated a higher insurance fee than what was observed in the data.

Two observations about Table B.2 stand out. First, over a quarter of the contract fees are estimated accurately down to the last cent (28.4 percent in the full sample). Allowing for ±5 percent (about ±\$6) error, two-fifths of the contracts have reasonably accurate price estimates.

Second, the error distributions are essentially the same across all subsamples. This fact is important for my argument that restricting the estimation sample to those people whose 2010 insurance premia are accurately calculated removes a large part of the measurement error in the switching benefit variable without biasing the estimation.

2.4 Robustness checks

I test the robustness of my results in several ways. First, I re-run all statistics and estimation procedures using different samples. As I described earlier, my baseline sample only contains drivers for whom the insurance premium calculation for 2010 was off by less than 5 percentage points. I argued that the small calculation error is a sign that there are no important unobserved price-relevant characteristics in this subsample, and therefore the estimated alternative prices and switching benefits are likely to be accurate as well.

I chose two additional subsamples for robustness testing. The first is a "1%" sample of potential switchers, meaning people for whom the calculated 2010 prices were within ±1 percent of the observed premia (and who have not dropped out of the sample before the first contract anniversary). The second dataset is a "50%" sample in the same vein. Whereas the baseline sample contains 6,766 observations in the *Campaign / No campaign* groups (i.e. excluding *Calendar* and *December* contracts), the narrower 1% sample has 4,246 data points and the wider 50% sample has 18,036.

10.14754/CEU.2015.04

Table B.2: Distribution of proportional insurance fee estimation errors in various subsamples

Percent.	All	Dropouts	Non-dropouts	Switchers	Non-switchers
5	-0.30	-0.30	-0.30	-0.30	-0.30
10	-0.22	-0.22	-0.22	-0.25	-0.22
15	-0.22	-0.21	-0.22	-0.22	-0.22
20	-0.17	-0.17	-0.17	-0.20	-0.17
25	-0.12	-0.10	-0.12	-0.17	-0.12
30	-0.07	-0.04	-0.07	-0.12	-0.07
35	-0.02	-0.02	-0.04	-0.07	-0.02
40	-0.02	-0.02	-0.02	-0.02	-0.02
45	-0.01	-0.00	-0.02	-0.02	-0.01
50	0.00	0.00	0.00	-0.00	0.00
55	0.00	0.00	0.00	0.00	0.00
60	0.00	0.00	0.00	0.00	0.00
65	0.00	0.00	0.00	0.00	0.00
70	0.03	0.00	0.04	0.00	0.06
75	0.12	0.11	0.12	0.11	0.13
80	0.15	0.15	0.15	0.15	0.15
85	0.18	0.18	0.18	0.18	0.18
90	0.20	0.18	0.21	0.18	0.23
95	0.33	0.33	0.33	0.32	0.33
Observations	24,187	4,805	19,382	4,261	15,121

10.14754/CEU.2015.04

In terms of basic descriptives, the alternative samples provide very similar results to Tables 2.1 and B.5. If anything, the *Campaign* and the *No campaign* drivers are even "closer" to one another in the 50% sample. Notably, their average insurance premia are practically indistinguishable.

There is one point where the 1% sample differs considerably from the 5% and the 50% samples: the market share of insurance companies. It seems that prices for Aegon customers are harder to calculate within a 1% accuracy band than for other companies, and therefore the market share of Aegon drops from about 45% in the 5% and 50% samples to about 20-25% in the 1% sample.

The main conclusion of Table 2.2, namely that *Campaign* and *No campaign* drivers have no differences in their price-relevant unobserved characteristics when all observables are controlled for, bears out in both the 1% and the 50% samples. Regarding unconditional dropout and switching rates, the three samples also produce practically identical results.

The logit estimates for binary marginal effects on contract switching are quantitatively similar in the baseline and the two alternative samples (see Table 2.5 in the main text, as well as Tables B.3 and B.4). In the 1% and 50% samples the campaign effect is even stronger: the conditional switching rate difference between *Campaign* and *No campaign* drivers amounts to 14-15 percentage points.

Another aspect in which the three samples differ slightly is the estimated coefficients on switching benefits. As the sample size increases, the estimates become smaller, which is most clearly seen by comparing Tables B.3 and B.4. This is likely a form of attenuation bias, since insurance prices (and therefore switching benefits) are estimated with larger errors in larger samples. The small difference between the 5% and the more accurate 1% sample suggests, on the other hand, that attenuation is not a serious issue in our baseline estimates.

Besides the sample-based robustness checks, I also ran the binary logit model by removing monetary switching benefits from the explanatory variables. The results are identical to the baseline estimates in Table 2.5, regardless of the control variables used in the estimation.

2.5 Additional tables

Table B.5 shows the share of insurers in the sample, broken down by treatment categories. The *Campaign* and the *No campaign* insurees are very similar to one another, with two caveats.

First, some insurers were relatively more successful in acquiring clients from one of the two groups in the entire sample (that is, without the ±5% restriction on price predictability). Specifically, KOBE and Posta did better in the *Campaign* group, and K&H and Waberer did worse (all the others were even). This difference might be due to advertising strategies, for example.[6]

Second, the sample restriction based on insurance fee predictability changed the market shares of insurers simply because some insurers' fee structures were more

[6]The case of Waberer is somewhat special. The company was forbidden by the regulator to acquire new clients between January 1st and April 15th in 2010 for failing to comply with administrative regulations. Hence its 0% market share in the *Campaign* group.

10.14754/CEU.2015.04

Table B.3: Logit marginal effects on the probability of switching at the first anniversary of the insurance contract

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Campaign	0.158***	0.148***	0.148***	0.146***	0.147***	0.140***	0.133***
	(0.043)	(0.043)	(0.044)	(0.044)	(0.043)	(0.044)	(0.043)
Switching benefit ($100)	0.087***	0.094***	0.095***	0.092***	0.099***	0.097***	0.082***
	(0.011)	(0.012)	(0.013)	(0.014)	(0.014)	(0.014)	(0.016)
Campaign × Sw. benefit	-0.006	0.000	-0.000	0.002	-0.001	0.006	-0.018
	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.052)	(0.050)
Risk controls		Yes	Yes	Yes	Yes	Yes	Yes
Personal controls			Yes	Yes	Yes	Yes	Yes
Vehicle controls				Yes	Yes	Yes	Yes
Payment controls					Yes	Yes	Yes
Contract channel controls						Yes	Yes
Insurer controls							Yes
Observations	4,246	4,246	4,246	4,246	4,246	4,246	4,246

Notes: All contracts start between January 2 and November 30, 2010.

10.14754/CEU.2015.04

Table B.4: Logit marginal effects on the probability of switching at the first anniversary of the insurance contract

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Campaign	0.156***	0.156***	0.155***	0.156***	0.158***	0.157***	0.154***
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
Switching benefit ($100)	0.075***	0.076***	0.073***	0.068***	0.069***	0.068***	0.078***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.008)
Campaign × Sw. benefit	0.013	0.012	0.015	0.014	0.012	0.013	0.010
	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)
Risk controls		Yes	Yes	Yes	Yes	Yes	Yes
Personal controls			Yes	Yes	Yes	Yes	Yes
Vehicle controls				Yes	Yes	Yes	Yes
Payment controls					Yes	Yes	Yes
Contract channel controls						Yes	Yes
Insurer controls							Yes
Observations	18,036	18,036	18,036	18,036	18,036	18,036	18,036

Notes: All contracts start between January 2 and November 30, 2010.

Table B.5: Market shares of insurance companies in the sample

	Campaign	No campaign	Mean diff.
Aegon	0.414	0.481	-0.067**
	(0.027)	(0.006)	(0.028)
Allianz	0.126	0.080	0.046**
	(0.018)	(0.003)	(0.019)
Generali	0.009	0.010	-0.000
	(0.005)	(0.001)	(0.005)
Genertel	0.156	0.148	0.009
	(0.020)	(0.004)	(0.021)
Groupama	0.074	0.064	0.010
	(0.014)	(0.003)	(0.015)
K&H	0.018	0.064	-0.046***
	(0.007)	(0.003)	(0.008)
KOBE	0.074	0.018	0.056***
	(0.014)	(0.002)	(0.015)
MKB	0.000	0.001	-0.001**
	(0.000)	(0.000)	(0.000)
Posta	0.031	0.008	0.023**
	(0.010)	(0.001)	(0.010)
Signal	0.000	0.004	-0.004***
	(0.000)	(0.001)	(0.001)
Union	0.089	0.093	-0.004
	(0.016)	(0.004)	(0.016)
Uniqa	0.009	0.009	0.001
	(0.005)	(0.001)	(0.005)
Waberer	0.000	0.020	-0.020***
	(0.000)	(0.002)	(0.002)
Observations	326	6,440	

predictable than others'. Although most of the market share changes were neutral across the *Campaign / No campaign* boundary, some were not. In particular, the market share differences for Aegon, Allianz, and MKB are an artifact of sample construction, rather than a sign of underlying differences in consumer choice.

10.14754/CEU.2015.04

Appendix C

Appendix for Chapter 3

10.14754/CEU.2015.04

3.1 Additional tables

Table C.1: Parameter estimates in the structural switching cost model with homogenous and heterogenous insurer fixed effects

	(1)	(2)
Intercept (switching cost)	0.653***	0.606***
	(0.111)	(0.110)
Campaign	-0.102	-0.163
	(0.250)	(0.228)
Intercept (θ)	0.967***	0.817***
	(0.082)	(0.099)
Capital resident (θ)	-0.150	-0.179*
	(0.095)	(0.101)
Online contracts (θ)	-0.581***	-0.649***
	(0.152)	(0.168)
Female (θ)	-0.194***	-0.176**
	(0.073)	(0.077)
Campaign (θ)	-0.905***	-0.898***
	(0.261)	(0.311)
ν_0	0.290***	0.179***
	(0.004)	(0.005)
ν_1	0.503***	0.455***
	(0.031)	(0.033)
Insurer fixed effects	Homogenous	Heterogenous
Observations	6,766	6,766

Table C.2: Parameter estimates in the structural switching cost model without insurer fixed effects

	(1)	(2)	(3)	(4)
Intercept (switching cost)	0.980***	1.086***	0.978***	1.012***
	(0.116)	(0.130)	(0.122)	(0.127)
Age		0.010***	0.010***	0.012***
		(0.002)	(0.002)	(0.003)
Online contracts				0.059
				(0.212)
Campaign	-0.063	-0.122	-0.116	-0.151
	(0.237)	(0.254)	(0.228)	(0.221)
Power (kW)			-0.010***	-0.009***
			(0.002)	(0.002)
Intercept (θ)	0.910***	0.844***	0.927***	0.897***
	(0.097)	(0.103)	(0.091)	(0.112)
Online contracts (θ)	-0.723***	-0.749***	-0.779***	-0.771**
	(0.168)	(0.175)	(0.170)	(0.301)
Campaign (θ)	-0.951***	-0.898***	-0.875***	-0.852***
	(0.336)	(0.343)	(0.272)	(0.269)
ν_0	0.491***	0.491***	0.491***	0.491***
	(0.007)	(0.007)	(0.007)	(0.007)
ν_1	0.452***	0.480***	0.473***	0.472***
	(0.033)	(0.037)	(0.036)	(0.036)
Risk and personal controls				Yes
Vehicle and contract controls				Yes
Personal controls (θ)	Yes	Yes	Yes	Yes
Observations	6,766	6,766	6,766	6,766

Table C.3: Parameter estimates in the structural switching cost model with all consumers (column 1) and recent entrants (columns 2-3) only

	(1)	(2)	(3)
Intercept (switching cost)	0.646***	0.388**	0.418
	(0.122)	(0.158)	(1.015)
Age	0.012***	0.032**	0.011***
	(0.003)	(0.015)	(0.004)
Online contracts	0.330	-0.223	0.365
	(0.234)	(0.541)	(0.251)
Campaign	-0.033	-0.184	-0.144
	(0.227)	(0.415)	(0.263)
Power (kW)	-0.010***	-0.010**	-0.010***
	(0.003)	(0.005)	(0.003)
Intercept (θ)	0.905***	0.989***	0.936***
	(0.096)	(0.115)	(0.093)
Online contracts (θ)	-0.876***	-0.515*	-0.816***
	(0.273)	(0.293)	(0.262)
Campaign (θ)	-0.975***	-0.832***	-0.913***
	(0.250)	(0.307)	(0.243)
ν_0	0.291***	0.293***	0.284***
	(0.004)	(0.006)	(0.005)
ν_1	0.526***	0.491***	0.536***
	(0.034)	(0.045)	(0.037)
Insurer fixed effects	Yes	Yes	Yes
Risk and personal controls	Yes	Yes	Yes
Vehicle and contract controls	Yes	Yes	Yes
Personal controls (θ)	Yes	Yes	Yes
Observations	6,766	3,032	5,500

10.14754/CEU.2015.04

Table C.4: Parameter estimates in the structural switching cost model with fully attentive consumers and common error variance over time

	(1)	(2)	(3)	(4)	(5)	(6)
Intercept (switching cost)	1.569***	1.573***	1.571***	1.552***	1.551***	1.552***
	(0.022)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Campaign	-0.399***	-0.401***	-0.397***	-0.396***	-0.393***	-0.392***
	(0.047)	(0.048)	(0.047)	(0.046)	(0.046)	(0.046)
ν	0.420***	0.423***	0.417***	0.410***	0.410***	0.410***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Insurer fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Risk controls		Yes	Yes	Yes	Yes	Yes
Personal controls			Yes	Yes	Yes	Yes
Vehicle controls				Yes	Yes	Yes
Payment controls					Yes	Yes
Contract channel controls						Yes
Observations	6,766	6,766	6,766	6,766	6,766	6,766