

# Saliency and Switching

András Kiss\*

JOB MARKET PAPER

December 29, 2014<sup>†</sup>

## Abstract

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 saliency of the switching opportunity for a subset of drivers. Using a micro-level dataset, I find that the campaign increases switching rates by 12 percentage points from a baseline of 20 percent. I also jointly estimate switching costs and consumer inattention in a structural model, showing that a quarter of insureds only consider switching because of the campaign, whereas almost half of them completely ignore the decision.

## 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

---

\*Department of Economics, Central European University, Budapest, Hungary. Email: [kiss.andras@ceu-budapest.edu](mailto:kiss.andras@ceu-budapest.edu). I am grateful to Péter Kondor, Miklós Koren, and Botond Kószegi for valuable guidance. I would also like to thank Paul Heidhues, Gábor Kézdi, Sergey Lychagin, José-Luis Moraga-González, Jenő Pál, Devin Pope, Ádám Szeidl, Péter Zsohár, and seminar participants at Central European University for useful comments and discussions, and especially dr. Gábor Vass and Árpád Antal for their support of the project and their help with data acquisition.

<sup>†</sup>The latest version of the paper is available at [www.andraskiss.com/jmp](http://www.andraskiss.com/jmp).

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.<sup>1</sup> 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 paper, 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 *saliency* 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.<sup>2</sup>

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

---

<sup>1</sup>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.

<sup>2</sup>\$50 is approximately the median saving in the sample, worth about one and a half days’ of average net wages.

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

In the first part of the paper, I measure the reduced-form effect of the campaign on switching decisions in a simple discrete choice model.<sup>3</sup> The dependent variable is whether a person switched contracts at the end of his

---

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

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 switching is suboptimal for many people.

For the second part of the empirical strategy, 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  $\theta$  (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.<sup>4</sup>

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.<sup>5</sup> 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.

A handful of recent empirical papers have tried to measure the sources of consumer inertia and studied the effect of salience in other choice situations. This is the first paper to identify the causal effect of choice salience on consumer switching in a natural experiment.

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. The main point of my paper, however, is the use of a natural experiment 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 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

---

<sup>4</sup>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.

<sup>5</sup>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.

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.<sup>6</sup>

The demand effects of the salience of taxes (Chetty et al. (2009), Finkelstein (2009)) and add-on prices (Hossain and Morgan (2006), Einav et al. (2014)), as well as the effect of limited investor attention on asset price adjustments (DellaVigna and Pollet (2009), Hirshleifer et al. (2009)) have been studied in recent empirical work. My paper complements this literature by identifying the effect of salience on consumer switching, a recurring decision problem with non-trivial economic consequences in many markets

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 paper, the bias could be corrected.

The rest of the paper 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 reduced-form specifications for contract switching. Section 5 contains the structural model and estimates, followed by the evaluation of two policy experiments, the discussion of the paper's 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.

---

<sup>6</sup>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.

## 2 The vehicle liability insurance market in Hungary

### 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.<sup>7</sup>

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.

#### 2.1.1 Contract switching prior to 2010

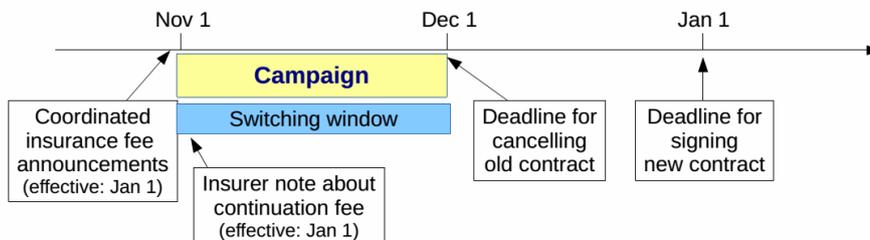
Until January 1, 2010 (in the *old regime*), contracts were required to coincide with the calendar year from the second year onwards. If, for example, someone bought a car in July 2008, his first insurance period only ran until December 31 of that year, and the second and subsequent periods covered the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

---

<sup>7</sup>The once-a-year price setting restriction has been removed from the regulation starting from January 2013. The focus of this paper is on years 2010 and 2011.

Figure 1: Time periods for switching were synchronized for all insurees 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.<sup>8</sup>

### 2.1.2 Contract switching since 2010

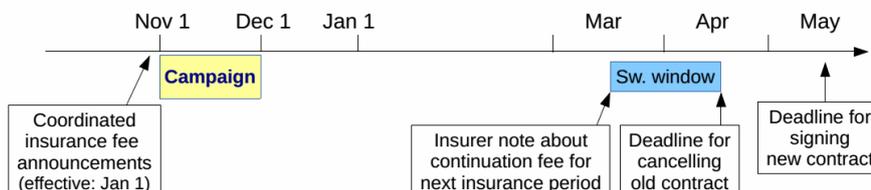
Following January 1st, 2010, all new insurance periods—including the first one—have become one year long. As a result, if someone bought a car in the middle of May 2010, the first time he could switch contracts was not in November 2010, but between mid-March and mid-April in 2011 (see Figure 2). The 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)

<sup>8</sup>The campaign had directly measurable effects on the salience of the switching opportunity in consumers' perception. A survey, commissioned by the Hungarian Competition 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)).

Figure 2: Time periods for switching are individual-specific since 2010



were still end-of-year switchers.<sup>9</sup>

## 2.2 Structure and calculation of the insurance premia

### 2.2.1 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.

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.

<sup>9</sup>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.

### 2.2.2 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 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.

### 2.2.3 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.<sup>10</sup> 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.<sup>11</sup>

---

<sup>10</sup>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.

<sup>11</sup>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

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.

### 3 Data and identification

I collected data from a mid-sized insurance broker company in Hungary for the empirical analysis.<sup>12</sup> 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.

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.

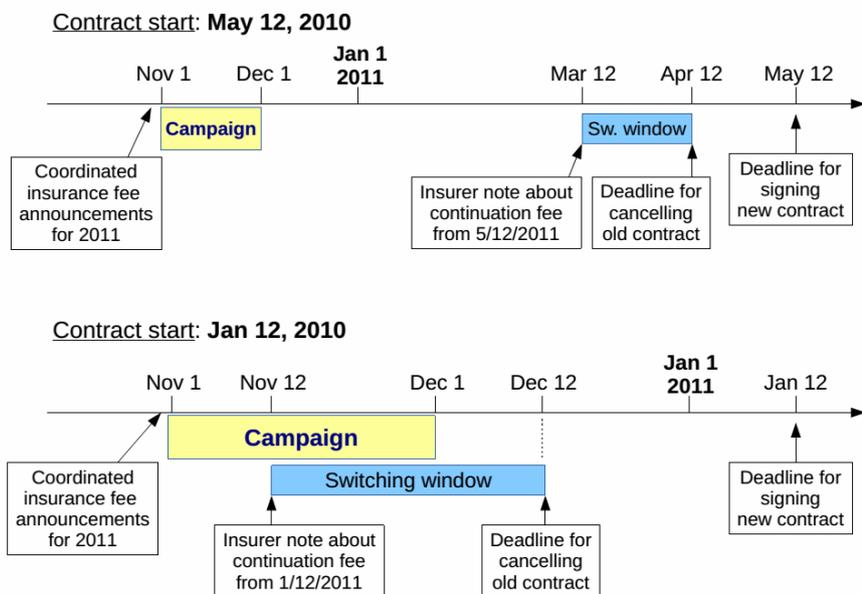
In addition, the company has a number of affiliates scattered across the country and one customer center in a large city, but both of these are rela-

---

the announcement itself to calculate prices.

<sup>12</sup>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 3: Identification of the campaign effect relies on the exogenous timing of individual switching periods



tively minor sources of new clients.

### 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 3 provides a detailed example.

The top panel of Figure 3 shows a contract with a starting date of May 12, 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).<sup>13</sup> 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.<sup>14</sup>

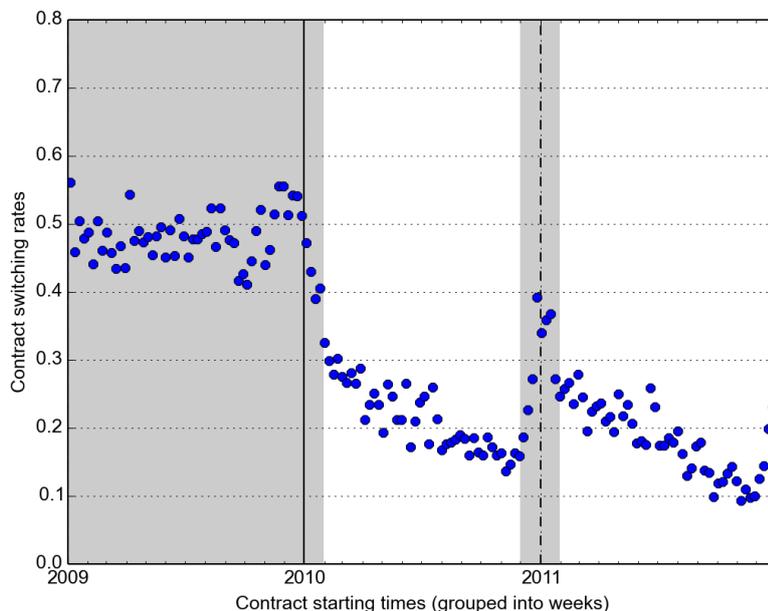
These objections to the identification argument can best be countered by examining Figure 4. In the figure, the horizontal axis is a timeline of contract starting times between January 1, 2009 and December 31, 2011. The vertical axis measures switching rates. Each dot in the graph represents the proportion of contracts that were cancelled at the first switching opportunity

---

<sup>13</sup>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.

<sup>14</sup>For example, there might be unobserved car discounts in January, which attract a larger proportion of financially savvy consumers. Differential macroeconomic effects within the year (e.g. changing economic outlook) might also play a role in the selection of car buyers.

Figure 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



within all the contracts started on the given week. I only included insurance contracts that were signed because of a vehicle acquisition.

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 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.

### 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,<sup>15</sup> 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).<sup>16</sup>

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

---

<sup>15</sup>The original data contains about 2.5 times as many observations as the baseline dataset.

<sup>16</sup>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.

contracts are therefore (with a slight rounding) the ones starting in January 2010. For the rest of the paper, I will use the terms *Campaign* and *No campaign* to denote the treatment and the control groups.<sup>17</sup>

Table 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.<sup>18</sup>

The main lesson from Table 1 is that the observable personal characteristics and 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.<sup>19</sup> 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.<sup>20</sup>

### 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

---

<sup>17</sup>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.

<sup>18</sup>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 paper.

<sup>19</sup>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-equipped gasoline-fuelled ones.

<sup>20</sup>Dollar amounts are calculated from Hungarian Forints using the 2010 average exchange rate of 208.15 HUF/USD.

Table 1: Descriptive statistics on insurees, vehicles, and contract characteristics

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

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 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.

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 1. Controlling 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.

Overall, the conclusion from Tables 1 and 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.

### 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.

Table 2: OLS estimates of insurance premium differences

	(1)	(2)	(3)	(4)	(5)
Campaign	0.077 (0.054)	0.115*** (0.038)	0.074** (0.031)	0.075** (0.031)	0.004 (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
<i>Observations</i>	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.

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.<sup>21</sup>

Table 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 *Campaign* 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. Still, the remaining *Campaign* insurees are more likely to change insurers than the remaining *No campaign*

<sup>21</sup>See the data appendix for more details about pinpointing switchers.

Table 3: Unconditional contract-ending hazard rates

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

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, could easily 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.

### 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

Table 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	
<i>Observations</i>	326	6,440	

they chosen any of the alternative offers. The price calculation methodology and the potential measurement issues are discussed in detail in the appendix.

Table 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 1), and is explained by the same peculiarities of sample construction I described earlier.

## 4 The campaign effect in reduced-form specifications

### 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, I use several sets of control variables. The results of the estimation are shown in Table 5.

The top panel of the table shows the estimated marginal effects for the

Table 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*** (0.031)	0.116*** (0.031)	0.116*** (0.031)	0.117*** (0.031)	0.120*** (0.031)	0.116*** (0.031)	0.119*** (0.031)
Switching benefit (\$100)	0.091*** (0.009)	0.096*** (0.009)	0.097*** (0.010)	0.091*** (0.011)	0.092*** (0.011)	0.091*** (0.011)	0.075*** (0.013)
Campaign $\times$ Sw. benefit	0.024 (0.038)	0.027 (0.038)	0.028 (0.038)	0.026 (0.038)	0.024 (0.038)	0.027 (0.038)	0.013 (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
<i>Observations</i>	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.

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.

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.<sup>22</sup>

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.<sup>23</sup>

Considering that switching benefits are below \$150 for almost everyone in the sample,<sup>24</sup> 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. I will deal with this issue later using a structural choice model.

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

---

<sup>22</sup>In a linear probability model, the specification in the first column of Table 5 yields a campaign coefficient of 0.172, exactly the same as the corresponding unconditional mean difference in Table 3. When control variables are also included, the linear probability estimates are essentially the same as those reported in Table 5.

<sup>23</sup>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.

<sup>24</sup>The 95th percentile of the benefit distribution is at \$148.

to gain less on average who need to be reminded by the campaign to shop around.

## 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.<sup>25</sup>

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 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.<sup>26</sup> I switched the reference group to January for presentation purposes. The results are shown in Table 14 in the appendix.

All coefficients are significant and stable across the different specifications.<sup>27</sup> 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.

## 5 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,

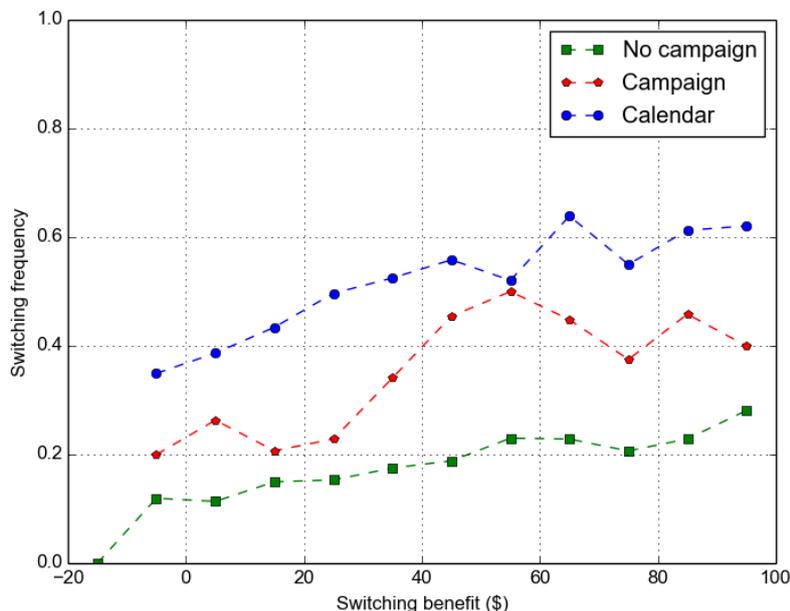
---

<sup>25</sup>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.

<sup>26</sup>I dropped the insignificant month-benefit interaction terms from the regression.

<sup>27</sup>I added the control groups in pairs to save space. The excluded columns are just like the ones that are included.

Figure 5: Switching rates as a function of financial savings



is whether people act as fully attentive consumers during the campaign period, or many of them remain ignorant of switching despite the campaign. In this part of the paper, 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 5 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.<sup>28</sup> 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).

<sup>28</sup>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.

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 5 also highlights 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.<sup>29</sup> 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  $\theta$  (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.

---

<sup>29</sup>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.

## 5.1 Model

### 5.1.1 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.<sup>30</sup>

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

$$U_{nj0} = \alpha X_{nj0} - p_{nj0} + \nu_0 \varepsilon_{nj0} \quad (1)$$

$p_{nj0}$  is the insurance premium that driver  $n$  would pay in contract  $j$  in 2010, and it affects utility negatively.  $X_{nj0}$  are insurer dummy variables, potentially interacted with individual characteristics; hence  $\alpha$  captures the insurer fixed effects. 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  $\varepsilon_{nj0}$ . For tractability, I restrict  $\varepsilon_{nj0}$  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  $\nu_0$ . 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:

$$P_{nj0} = \frac{\exp\left(\frac{\alpha}{\nu_0} X_{nj0} - \frac{1}{\nu_0} p_{nj0}\right)}{\sum_{i \in J^0} \exp\left(\frac{\alpha}{\nu_0} X_{ni0} - \frac{1}{\nu_0} p_{ni0}\right)} \quad (2)$$

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

### 5.1.2 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:

$$U_{nj1} = \alpha X_{nj1} - p_{nj1} - \beta Z_{nj1} + \nu_1 \varepsilon_{nj1} \quad (3)$$

---

<sup>30</sup>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 5 to the model.

The insurer fixed effects are time-invariant.<sup>31</sup> The price variable  $p_{nj1}$  reflects what the consumer would have to pay in period 1 for contract  $j$ , while the random term  $\varepsilon_{nj1}$  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.

$Z_{nj1}$  contains the determinants of switching costs, interacted with a switching indicator. In the simplest case,  $Z_{nj1} = 1$ , and the corresponding  $\beta$  parameter (the intercept) measures the average costs of switching across the sample. The period 1 choice probabilities are analogous to expression (2).

### 5.1.3 Inattention to switching

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

$$j_n^1 = \begin{cases} \operatorname{argmax}_j U_{nj1} & \text{with prob. } \theta_n \\ j_n^0 & \text{with prob. } 1 - \theta_n \end{cases} \quad (4)$$

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

$$\theta_n = \frac{1}{1 + \exp(\gamma W_n)} \quad (5)$$

Thus, a consumer takes a conscious utility-maximizing choice (allowing for switching costs) in period 1 with probability  $\theta_n$ , 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.

### 5.1.4 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 6 provides a stylized graphical explanation.

---

<sup>31</sup>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).

The horizontal and vertical axes in Figure 6 measure the financial savings and the probability of switching, similarly to Figure 5. 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 6 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 5 and 6, whereas the switching cost equals the amount of financial savings for which the observed switching rates are equal to those predicted by the model.

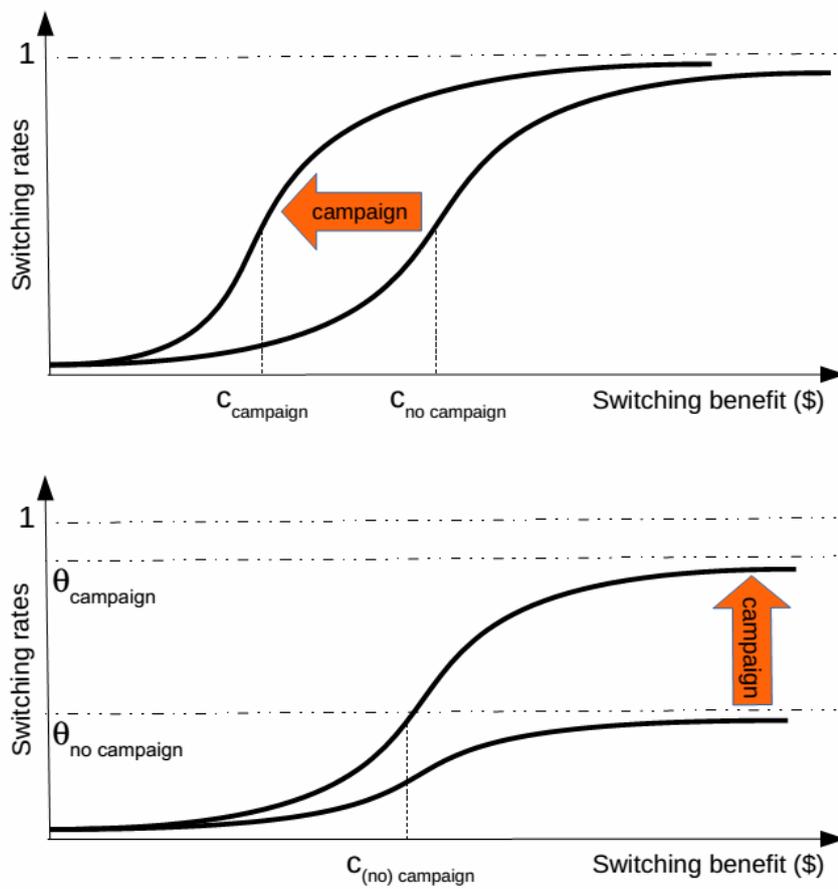
### 5.1.5 Censoring

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,<sup>32</sup> 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.

---

<sup>32</sup>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 6: Structural effects on switching costs vs attention



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.

### 5.1.6 Aggregate choice probabilities

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

$$P_n = P_{nj^0} \cdot P_{nj^1} \quad (6)$$

where  $P_{nj^0}$  is given by equation (2) for the actual choice  $j = j^0$ , and

$$P_{nj^1} = \theta_n \cdot \frac{\exp\left(\frac{\alpha}{\nu_1} X_{nj^1} - \frac{1}{\nu_1} p_{nj^1} - \frac{\beta}{\nu_1} Z_{nj^1}\right)}{\sum_{i \in J^1} \exp\left(\frac{\alpha}{\nu_1} X_{ni1} - \frac{1}{\nu_1} p_{ni1} - \frac{\beta}{\nu_1} Z_{ni1}\right)} + (1 - \theta_n) \cdot d_n \quad (7)$$

when an observation in period 1 is not censored, and

$$P_{nj^1} = \theta_n \cdot \left[ 1 - \frac{\exp\left(\frac{\alpha}{\nu_1} X_{nj^0} - \frac{1}{\nu_1} p_{nj^0} - \frac{\beta}{\nu_1} Z_{nj^0}\right)}{\sum_{i \in J^1} \exp\left(\frac{\alpha}{\nu_1} X_{ni1} - \frac{1}{\nu_1} p_{ni1} - \frac{\beta}{\nu_1} Z_{ni1}\right)} \right] \quad (8)$$

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.

Taking the logarithm of expression (6) and summing over the individuals yields the objective function of a maximum likelihood estimation procedure. The log-likelihood function has a closed form and can be maximized using standard optimization procedures.

## 5.2 Structural estimation results

Table 6 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  $\beta$  in equation 3 and measured in \$100), the second panel shows the estimated parameters influencing inattention ( $\gamma$ ), the third panel marks the applied sets of control variables, followed by the estimates of the random scale parameters ( $\nu$ ). 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 in the \$60-80 range, which is a plausible outcome given industry reports in trade press articles.<sup>33</sup> 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 7 shows the values of  $\theta$  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 7.

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 7). Another 12 percent said that they took care of switching only because of the constant reminders,

---

<sup>33</sup>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.

Table 6: Parameter estimates in the structural switching cost model with inattentive consumers

	(1)	(2)	(3)	(4)
Intercept (switching cost)	0.653*** (0.111)	0.695*** (0.116)	0.627*** (0.121)	0.646*** (0.122)
Age		0.013*** (0.003)	0.013*** (0.003)	0.012*** (0.003)
Online contracts				0.330 (0.234)
Campaign	-0.102 (0.250)	-0.105 (0.256)	-0.068 (0.240)	-0.033 (0.227)
Power (kW)			-0.010*** (0.003)	-0.010*** (0.003)
Intercept ( $\theta$ )	0.967*** (0.082)	0.940*** (0.084)	0.986*** (0.078)	0.905*** (0.096)
Online contracts ( $\theta$ )	-0.581*** (0.152)	-0.596*** (0.154)	-0.644*** (0.155)	-0.876*** (0.273)
Campaign ( $\theta$ )	-0.905*** (0.261)	-0.906*** (0.271)	-0.917*** (0.241)	-0.975*** (0.250)
$\nu_0$	0.290*** (0.004)	0.290*** (0.004)	0.291*** (0.004)	0.291*** (0.004)
$\nu_1$	0.503*** (0.031)	0.516*** (0.033)	0.525*** (0.034)	0.526*** (0.034)
Insurer fixed effects	Yes	Yes	Yes	Yes
Risk and personal controls				Yes
Vehicle and contract controls				Yes
Personal controls ( $\theta$ )	Yes	Yes	Yes	Yes
<i>Observations</i>	6,766	6,766	6,766	6,766

Table 7: 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%

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 7, where the weights correspond to the share of people using online contracting channels in the entire population.

## 5.3 Robustness checks

### 5.3.1 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  $\alpha$  parameters to take on individual-specific (but time-invariant) values from independent normal distributions.<sup>34</sup> The results are shown in Table 15 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.

### 5.3.2 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

---

<sup>34</sup>That is, instead of estimating the  $\alpha$  coefficients, I estimate the mean and the variance of their (independent normal) distributions.

(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.

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

Allowing for heterogeneous 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 16 in the appendix. Column 1 of the table shows the baseline estimates from Table 6 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.

#### 5.3.4 Full attention

We can confirm the benefit of explicitly modeling inattention by running an estimation in which  $\theta$  is set to 1; that is, all consumers are assumed to be fully attentive. The result of this exercise is shown in Table 17 in the appendix. Switching costs increase from \$65 to about \$550, which is an unreasonably high value.<sup>35</sup> 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.

#### 5.3.5 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 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 8 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 6), the baseline attention level, and the effect of the media campaign on the share of attentive people.

The lesson from Table 8 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

---

<sup>35</sup>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.

Table 8: Structural estimates are robust to a wide range of price calculation errors

	Error threshold in 2010 contract price calculation						
	$\pm 1\%$	$\pm 5\%$	$\pm 10\%$	$\pm 15\%$	$\pm 20\%$	$\pm 30\%$	$\pm 50\%$
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%
<i>Observations</i>	4,246	6,766	7,827	9,749	13,219	16,140	18,036

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.

## 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 6, 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 proce-

ture as Hortacısu 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:

$$\Delta E(CS)_n = \nu_1 \left[ \log \left( \sum_{j \in J_n^{1,CF}} e^{\frac{1}{\nu_1} V_{nj1}^{CF}} \right) - \log \left( \sum_{j \in J_n^{1,BL}} e^{\frac{1}{\nu_1} V_{nj1}^{BL}} \right) \right] \quad (9)$$

where  $BL$  stands for the baseline scenario and  $CF$  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)).

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  $\theta_n$ , 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 (9) to:

$$\Delta E(CS)_n = V_{nj1}^{CF} - V_{nj1}^{BL} \quad (10)$$

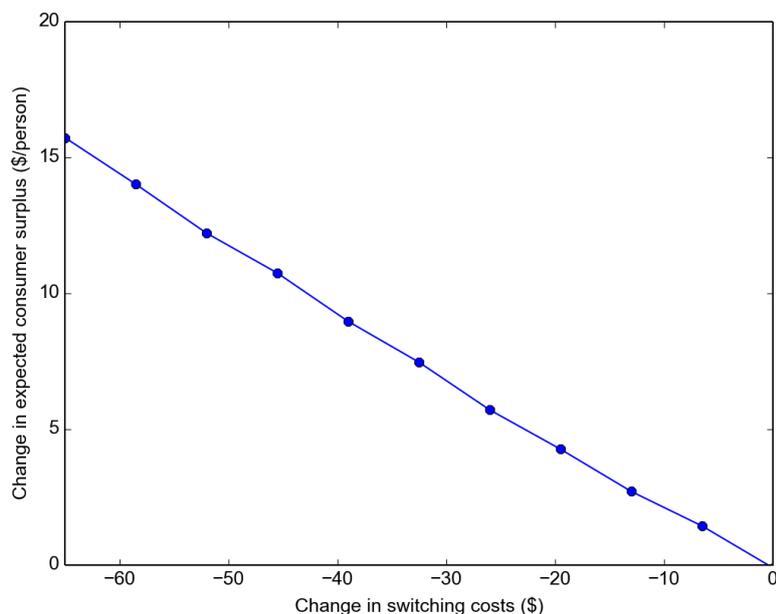
Since utility is measured in dollar terms, the interpretation of the values for  $\Delta E(CS)_n$  is also in dollars.

The first hypothetical policy intervention (switching cost reduction) acts by reducing the intercept of  $\beta Z_{nj1}$  in equation (3), hence it will make  $V_{nj1}^{CF}$  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  $V_{nj1}^{CF} = V_{nj1}^{BL}$  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),  $\Delta E(CS)_n$  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  $(\theta_1, \dots, \theta_N)$  and averaging the result.

Figure 7: Reducing switching costs by \$65 (to \$0) raises consumer surplus by \$15.6/person

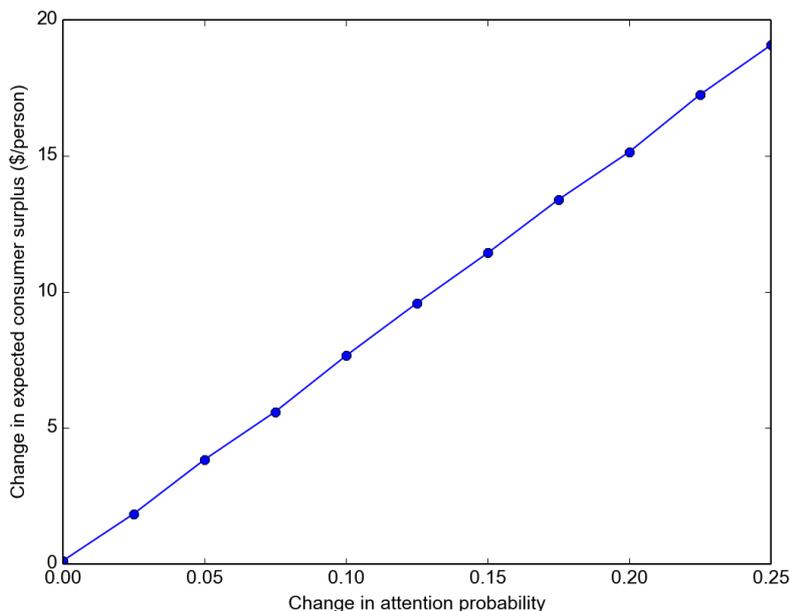


Figures 7 and 8 show the results of the counterfactual simulations. In Figure 7, I vary the cost of switching between \$0 and \$65 on the horizontal axis, and plot the average change in expected consumer surplus on the vertical axis. As the figure shows, the effect is close to linear in the range of costs that I consider. One 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 8, 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

Figure 8: A 10 percentage point increase in attention probability raises consumer surplus by \$7.7/person



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.

## 7 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.<sup>36</sup> 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

---

<sup>36</sup>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](http://www.andraskiss.com/research).

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 paper'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 choose 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 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 costs 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.<sup>37</sup> However, there are also several counterarguments that paint a less optimistic picture. First, many insureds 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

---

<sup>37</sup>Unfortunately, my dataset is too short to test for learning.

are much higher. The likely implication is that people learn slowly, if at all, from errors of omission.

The results of the paper 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.<sup>38</sup> 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.<sup>39</sup> 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.

## 8 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.

I also 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

---

<sup>38</sup>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.

<sup>39</sup>Under current rules, the first installment is only due three months after the switching deadline.

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.

The results of the paper 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.

# A Data appendix

## A.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 birth-dates, 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.<sup>40</sup>

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.

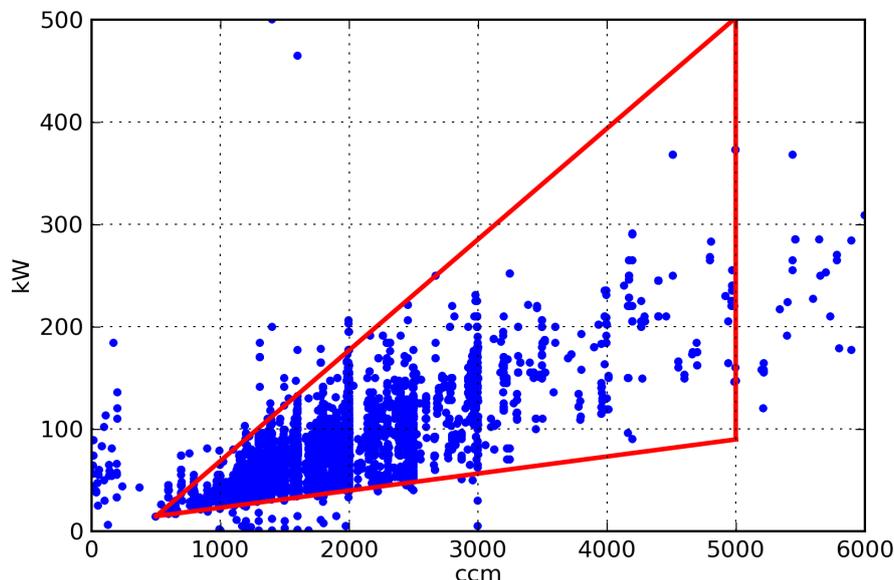
I also excluded contracts in the M1–M4 risk categories. These drivers have been found at fault in recent accidents. Since I have no information on the date of the damage claims, the estimation of the alternative insurance premia would contain considerable inaccuracy, and hence worsen the quality of my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

---

<sup>40</sup>The number of simultaneously insured vehicles by the same person (sometimes over half a dozen) points towards a gaming of the system.

Figure 9: Valid and excluded vehicle engine characteristics



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 9 plots the volume-power relationship seen in the data. I chose to exclude all records that fell outside the red triangle.<sup>41</sup>

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 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 character-

---

<sup>41</sup>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).

istics 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 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  $\pm 1\%$ , one with  $\pm 5\%$ , and one with  $\pm 50\%$  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.

## A.2 Establishing contract switching

In the insurance broker’s database, contracts referring 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 occurred.

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.<sup>42</sup> Figure 10 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.<sup>43</sup> 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.

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 10 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.

### A.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.

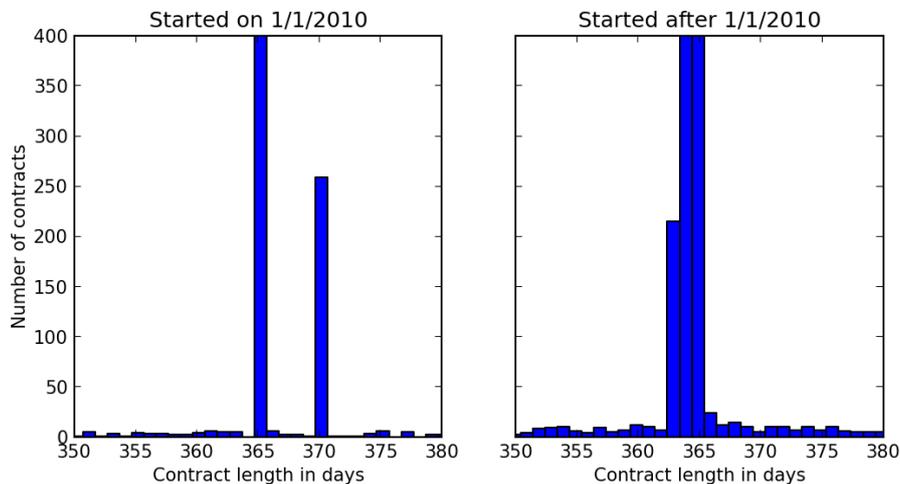
To estimate switching costs from observed choices, I need a good predictor of switching benefits. I use the publicly available pricing tables of insurance companies, together with relevant driver- and car-specific information, to

---

<sup>42</sup>I allow for some leeway in the recording of the termination date.

<sup>43</sup>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 10: Contract length distributions around 365 days for contracts started in 2010



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 9 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.

Table 9: 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

### A.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.<sup>44</sup> 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 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

<sup>44</sup>Considering that the calculations involved historic insurance fees, I had no opportunity to use currently operating price calculators.

the price calculation procedure.

### **A.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 10 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 10 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  $\pm 5$  percent (about  $\pm \$6$ ) 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.

## **B Robustness checks for the reduced-form results**

I test the robustness of my reduced-form results in two 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  $\pm 1$  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

Table 10: 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
<i>Observations</i>	24,187	4,805	19,382	4,261	15,121

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.

In terms of basic descriptives, the alternative samples provide very similar results to Tables 1 and 13. 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, 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 5 in the main text, as well as Tables 11 and 12). 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 11 and 12. 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 5, regardless of the control variables used in the estimation.

Table 11: 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*** (0.043)	0.148*** (0.043)	0.148*** (0.044)	0.146*** (0.044)	0.147*** (0.043)	0.140*** (0.044)	0.133*** (0.043)
Switching benefit (\$100)	0.087*** (0.011)	0.094*** (0.012)	0.095*** (0.013)	0.092*** (0.014)	0.099*** (0.014)	0.097*** (0.014)	0.082*** (0.016)
Campaign $\times$ Sw. benefit	-0.006 (0.051)	0.000 (0.051)	-0.000 (0.051)	0.002 (0.051)	-0.001 (0.051)	0.006 (0.052)	-0.018 (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
<i>Observations</i>	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.

Table 12: 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*** (0.017)	0.156*** (0.017)	0.155*** (0.017)	0.156*** (0.017)	0.158*** (0.017)	0.157*** (0.017)	0.154*** (0.017)
Switching benefit (\$100)	0.075*** (0.006)	0.076*** (0.006)	0.073*** (0.006)	0.068*** (0.006)	0.069*** (0.006)	0.068*** (0.006)	0.078*** (0.008)
Campaign $\times$ Sw. benefit	0.013 (0.022)	0.012 (0.022)	0.015 (0.022)	0.014 (0.022)	0.012 (0.022)	0.013 (0.022)	0.010 (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
<i>Observations</i>	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.

## C Additional tables

### C.1 Insurer market shares

Table 13 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  $\pm 5\%$  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.<sup>45</sup>

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

### C.2 Regression results

Tables 14-17 show additional regression results that are referenced in the main text.

---

<sup>45</sup>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.

Table 13: Market shares of insurance companies in the sample

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

Table 14: Logit marginal effects on the probability of switching at the first anniversary of the insurance contract

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

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

Table 15: Parameter estimates in the structural switching cost model with homogenous and heterogenous insurer fixed effects

	(1)	(2)
Intercept (switching cost)	0.653*** (0.111)	0.606*** (0.110)
Campaign	-0.102 (0.250)	-0.163 (0.228)
Intercept ( $\theta$ )	0.967*** (0.082)	0.817*** (0.099)
Capital resident ( $\theta$ )	-0.150 (0.095)	-0.179* (0.101)
Online contracts ( $\theta$ )	-0.581*** (0.152)	-0.649*** (0.168)
Female ( $\theta$ )	-0.194*** (0.073)	-0.176** (0.077)
Campaign ( $\theta$ )	-0.905*** (0.261)	-0.898*** (0.311)
$\nu_0$	0.290*** (0.004)	0.179*** (0.005)
$\nu_1$	0.503*** (0.031)	0.455*** (0.033)
Insurer fixed effects	Homogenous	Heterogenous
<i>Observations</i>	6,766	6,766

Table 16: 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*** (0.122)	0.388** (0.158)	0.418 (1.015)
Age	0.012*** (0.003)	0.032** (0.015)	0.011*** (0.004)
Online contracts	0.330 (0.234)	-0.223 (0.541)	0.365 (0.251)
Campaign	-0.033 (0.227)	-0.184 (0.415)	-0.144 (0.263)
Power (kW)	-0.010*** (0.003)	-0.010** (0.005)	-0.010*** (0.003)
Intercept ( $\theta$ )	0.905*** (0.096)	0.989*** (0.115)	0.936*** (0.093)
Online contracts ( $\theta$ )	-0.876*** (0.273)	-0.515* (0.293)	-0.816*** (0.262)
Campaign ( $\theta$ )	-0.975*** (0.250)	-0.832*** (0.307)	-0.913*** (0.243)
$\nu_0$	0.291*** (0.004)	0.293*** (0.006)	0.284*** (0.005)
$\nu_1$	0.526*** (0.034)	0.491*** (0.045)	0.536*** (0.037)
Insurer fixed effects	Yes	Yes	Yes
Risk and personal controls	Yes	Yes	Yes
Vehicle and contract controls	Yes	Yes	Yes
Personal controls ( $\theta$ )	Yes	Yes	Yes
<i>Observations</i>	6,766	3,032	5,500

Table 17: Parameter estimates in the structural switching cost model with fully attentive consumers

	(1)	(2)	(3)	(4)	(5)	(6)
Intercept (switching cost)	5.245*** (0.352)	5.240*** (0.351)	5.348*** (0.373)	5.532*** (0.409)	5.602*** (0.421)	5.699*** (0.439)
Campaign	-1.183*** (0.186)	-1.184*** (0.186)	-1.221*** (0.193)	-1.258*** (0.202)	-1.282*** (0.205)	-1.294*** (0.209)
$\nu_0$	0.293*** (0.005)	0.293*** (0.005)	0.293*** (0.005)	0.293*** (0.005)	0.293*** (0.005)	0.293*** (0.005)
$\nu_1$	1.416*** (0.090)	1.416*** (0.090)	1.446*** (0.097)	1.494*** (0.106)	1.510*** (0.109)	1.534*** (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
<i>Observations</i>	6,766	6,766	6,766	6,766	6,766	6,766

## References

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