NHH

INSTITUTT FOR SAMFUNNSØKONOMI

DEPARTMENT OF ECONOMICS

SAM 23 2018

ISSN: 0804-6824

November 2018

Discussion paper

The Poking Effect: Price Changes, Information, and Inertia in the Market for Mobile Subscriptions

BY

Bjørn-Atle Reme, Helene Lie Røhr AND **Morten Sæthre**

Norges Handelshøyskole

NORWEGIAN SCHOOL OF ECONOMICS

The Poking Effect: Price Changes, Information, and Inertia in the Market for Mobile Subscriptions*

Bjørn-Atle Reme† Helene Lie Røhr‡ Morten Sæthre§

November, 2018

Abstract

We study consumer inertia in the mobile subscription market, focusing on the decision of whether to switch to a competing provider. To identify the extent of inertia, we exploit price changes faced by 270,000 consumers of a large telecom provider. We document that the propensity to switch provider after the price change increases among consumers whose costs decrease with the new prices. Furthermore, we find that the increase is largest right after consumers are informed of the upcoming change—during the two months prior to the tariff change—as opposed to when the price change is implemented. From these findings, we infer what we call a *poking effect*; the information of an upcoming price change causes consumers to engage in searches for alternative offers, leading to increased switching. We supplement the analysis with a survey and find indications that the poking effect is due to consumer inattention. To separate the effect on attention from the reaction to the actual price change, we estimate a model of consumer choice with limited attention. We find that when consumers are poked, it increases the share of consumers becoming attentive to competing offers. This leads many consumers to switch providers earlier than they would otherwise, explaining why they leave even though their terms with their current company improve.

Keywords: Consumer Inertia, Inattention, Consumer Behavior, Telecom, Churn.

---

*We wish to thank Eva Ascarza, Matthew Backus, Heléne Berg, Alexander W. Cappelen, Pierre Dubois, Jon Fiva, Thomas de Haan, Andreas Kotsadam, Johanna Mollerstrom and Katja Seim for their discussions and invaluable inputs. The paper has also benefited from the comments of seminar and conference participants at VATT Institute for Economic Research, the Norwegian School of Economics, the Frisch Centre, the Norwegian Business School, EARIE, ESA, IIOC, NCBEE, and BECCLE Competition Policy Conference.

†Telenor Research and The Choice Lab. E-mail: Bjorn-Atle.Reme@telenor.com. The author's time is funded by Telenor Research. The views expressed in this paper are those of the author, and do not necessarily reflect the views of Telenor Research.

‡Telenor Research and BI Norwegian Business School. E-mail: Helene-Lie.Rohr@telenor.com. The author's time is funded by Telenor Research and the Research Council of Norway. The views expressed in this paper are those of the author, and do not necessarily reflect the views of Telenor Research.

§Norwegian School of Economics. E-mail: morten.saethre@nhh.no

1 Introduction

Consumer inertia is of interest to both business managers and policy makers, as it can have significant effects on market outcomes and dynamics. While previous studies have documented the existence and effects of choice inertia,[1] our paper contributes to the literature by providing evidence on the role of limited attention. Utilizing a text message notification sent from a telecom provider to approximately 270,000 consumers regarding an upcoming price change, we present evidence of what we refer to as the "poking effect": the propensity to switch provider (also referred to as "churn") increases following the notification, not only among consumers whose costs increase from the price change, but also among consumers whose costs decrease. The estimated effect implies a non-negligible increase in the churn rate in the range of 50 to 80 percent for consumers whose costs decrease. We suggest that the main mechanism behind the poking effect is a change in consumer attention from the price change which triggers search. This is supported by survey evidence from 451 consumers.

Even though we can observe the consequence of increased attention, the economic significance of attention will depend on the preferences of consumers and configuration of options in the market. Therefore, we estimate a choice model that allows us to quantify both the share of consumers who are attentive (interpreted as active choice behavior) and the effect of the poke on attention.[2] The model describes consumers' decisions to churn as a two-step process. First, consumers become attentive or not. Second, conditional on being attentive, the consumers choose their price plan (and provider). An important feature of the choice model is that it allows us to separate the effect of the poke on consumers' attention from that of removing their current plans.[3] We estimate that the probability of being attentive in a given month is roughly five percent in the period prior to notification, a figure that is very close to self-reported frequency of price search from our survey. Furthermore, our estimates imply that the information regarding the price change more than doubled the share of consumers being attentive.

Taken together our results indicate that many consumers are passive buyers in the sense that their choice of provider is "closed"—they do not reoptimize unless some salient event triggers them, for instance a price change. This type of behavioral pattern is consistent with inattention interacting with search costs, where, for instance, salient events like aggressive advertising campaigns signal that market prices have changed significantly, and that it is worthwhile to engage in searches for alternative offers. Hence, the price change catches the limited attention of consumers and provides a rationale for considering alternative providers.[4]

[1] See for example Madrian and Shea (2001); Handel (2013); Marzilli Ericson (2014); Miravete and Palacios-Huerta (2014a); Grubb (2015b); Ho et al. (2017)

[2] Our model is similar to that estimated by Hortaçsu, Madanizadeh, and Puller (2017), though our data are more informative about changes in attention due to variation coming from the removal of the currently chosen plan for one group of consumers, together with the information (SMS notification) these consumers received prior to the removal.

[3] Forcing consumers to migrate from their current price plan to a plan chosen by the operator may increase consumers' propensity to churn because of both the removal of their chosen plan (even if attention is unaffected) and the increase in attention to their choice of price plan when they are informed about the change.

[4] Note that search costs without inattention is not sufficient for explaining this behavioral pattern since the poking effect is only occurring among the consumers affected by the price change. If search costs alone was the

We provide two further pieces of evidence consistent with inattention (or, rather, an increase in attention) being the main mechanism behind the observed increase in switching. First, the survey evidence shows that consumers of the relevant operator rarely check competing offers, i.e., they are inattentive. At the same time, a substantial share of respondents say they would check competing offers if their provider lowered the price on their plan. Hence, it is credible that a price change triggers search. Secondly, to shed light on the potential role of choice mistakes, we use information on price plans offered by competing providers, and show that among the better off consumers who churn, 80 percent could do even better with their newly chosen provider. Hence, these consumers generally seem able to make cost-minimizing choices in this market.

Our dataset has two features that make it particularly suited for studying consumer inertia. First, the change in monthly payments from being migrated to a new price plan varies substantially among consumers, where some gain and some lose, depending on their previous price plan and usage profile. The consumers who gain are particularly interesting for our analysis in terms of identifying the poking effect. Second, consumers are informed about the change two months prior to the (forced) migration, while none of the competitors change their prices or offered plans during the period. This feature allows us to study the isolated effect of the information on switching, as opposed to experience with new prices. Given the rapid and strong consumer response following the notification, i.e., before the price change is implemented, we can conclude that the poking effect is triggered by the information on the upcoming change, and results in consumers revising their choice of provider.

Our paper is related to several recent studies focusing on the different aspects of choice inertia.[5] Consistent with our findings, Ascarza, Iyengar, and Schleicher (2016) show an increase in the churn rate among consumers exposed to a service provider's retention program, and propose reduced inertia as an important explanation for why the retention program failed. The literature has also been concerned with understanding the causes of inertia, since such knowledge has important implications for both the need for, and design of, policies aimed at increasing consumer switching. In this regard, an important question is the extent to which inertia is driven by frictions in consumer choice (such as search costs, switching costs, or inattention), or whether low rates of switching from current choices primarily result from (potentially unobserved) preference heterogeneity.[6] If preference heterogeneity is the most important driver of consumer choice, apparent inertia (in the form of low switching rates) can result from consumers being well matched with their current choice. Our finding of increased churn among consumers who would be better off under the new contract terms is hard to reconcile with preference heterogeneity alone driving differences in choices. This points to an important role of choice frictions in this market, where we argue that the nature of the changes we document

---

mechanism, we should expect some effect among subscribers unaffected by the price change, but with a similar usage profile as the affected customers. This we do not find in our data.

[5]See Madrian and Shea (2001), Handel (2013), Miravete and Palacios-Huerta (2014a),Marzilli Ericson (2014), Grubb (2015b), and Ho, Hogan, and Morton (2017). Also related are the literatures on the status quo bias (Samuelson and Zeckhauser, 1988), the effect of reminders (Karlan et al., 2016; Hanna et al., 2014) and salience (see, for instance, Bordalo et al., 2013).

[6]See, e.g., Hortaçsu, Madanizadeh, and Puller (2017) and Dube, Hitsch, and Rossi (2010).

in consumer choices is most readily explained by (in)attention being the driving force.

Our paper proceeds as follows: Section 2 describes the data and the market environment. Section 3 entails the reduced form difference-in-difference estimation of the poking effect and survey results. In Section 4, we present and estimate the structural model of consumer attention and choice. Conclusions are drawn in Section 5.

2 Data

Our study is based on data from a European mobile operator covering about 1.1 million consumers, which is the universe of the operator's postpaid consumers. We have up to 23 months of billing data for each consumer, from September 2013 to July 2015.[7] In addition to monthly observations on mobile usage, including calling, SMS, and mobile data, the dataset contains consumer characteristics such as gender, age, mobile spending, price plan, sign-up date for current price plan, and churn.

In an effort to simplify its subscription portfolio, the mobile operator removed 13 calling plans in May 2014, affecting about 270,000 consumers.[8] These consumers were moved to one of two predefined plans depending on the specific plan of the consumer prior to removal.[9] An important feature of our study is that the consumers were notified about the upcoming change via SMS in March, i.e., 2 months prior to when it was effectuated.

Our data cover 6 months prior to the SMS notification, the month of notification, and 16 months after notification (14 months after reassignment). A central outcome of our analysis is how the price change affected consumers' propensity to churn. When consumers churn, we no longer have data on their mobile usage or spending. Hence, in each month, we only have information on consumers of the mobile operator who have chosen not to churn.

The reassignment to new price plans implied a major change in the tariff structure for the affected consumers. The assigned calling plans featured three-part tariffs—i.e., a monthly fee, an included allowance, and a marginal price above the allowance—while most of the previous plans were either two-part or simple linear tariffs. Consumers on 7 of the 13 removed price plans were assigned to a price plan that included 100 minutes of calling, whereas consumers of the remaining 6 price plans were assigned to a separate, new plan that included 100 SMS messages.[10] For the most part, these price plans were not sold simultaneously. Hence, we see some plans strictly dominating others. With one exception, none of the removed plans had

[7] We only include consumers who carry the cost of their own mobile usage; hence, very young (age below 21 years) and business customers are excluded. After these restrictions, we are left with approximately 1 million consumers in our final dataset.

[8] In the first month of our sample, we have about 320,000 consumers on plans that would be removed, though this number decreased to 270,000 in the month prior to notification, due to both churn and consumers switching to other calling plans offered by the same provider.

[9] One of these plans had already existed on the market for 3 years, whereas the other was introduced at the same time as the plan removal. Neither of the price plans were actively marketed prior to March 2014.

[10] Table A.1 and Table A.2 in the Appendix give an overview of the important prices for the removed plans along with the new, assigned plans.

new consumers during our sample period.[11] The new price plans differed from the old ones both in terms of fixed fees, marginal prices, included volumes, and name. If the consumers did not make an active choice, such as changing to another provider or another plan within the company, the new plan would be effective from May.

For consumers who churn, we know the identity of their new provider, though we do not know which price plan they chose from the new provider's portfolio. However, using data on the price portfolio of the competitors for the relevant time window (cf. Table A.3), we can check if a churning consumer could leave for products that are less expensive, and how much less expensive that product would be (see Figure 4 and the accompanying discussion).[12] Information on competing providers' prices is also used when we estimate the choice model in Section 4.

2.1 Survey data

We conducted a survey with 451 postpaid customers of the telecom company. The survey was administered by Kantar TNS, a market research agency with presence in over 80 countries.[13] In addition to the survey questions (see section 3.3 for a detailed presentation), we collected demographics, such as gender, education, age and income. Descriptive statistics of the respondents along the demographics can be found in Table A.5.

2.2 Identification strategy

The main purpose of our analysis is to investigate the existence and severity of the poking effect—the churn caused by the information about the price change, as opposed to churn among consumers experiencing an increase in expenditure due to the price change. In our data, 270,000 consumers were notified by SMS about the upcoming price change, whom we denote as *poked*, whereas the remaining consumers were not poked (received neither an SMS nor a price change).[14] The price change had a heterogeneous effect on the costs for the poked consumers, depending on their prior plan and usage patterns. To identify the poking effect, we divide consumers in groups based on how the new prices would influence expected costs, where the groups are defined by (non-overlapping) ranges of changes in costs. We estimate the effect of notifying the consumers on churn using a difference-in-differences (DiD) estimator for each group of consumers, with poked consumers as the treatment group and non-poked consumers as the control group. The goal of this exercise is to isolate the reaction to the notification for consumers whose costs would be unchanged (or even decrease) following the

[11] One of the price plans sees an influx of 6,200 new consumers from the beginning of our sample to notification, which is small compared with both the total sample and the 100,000 consumers who were initially on the same plan. This group is included in the sample for the results presented below, though our results are unchanged if these consumers are excluded from the sample.

[12] Note that for technical data warehouse reasons, this information is only available for approximately one-half of the relevant subgroup of consumers.

[13] See http://www.tnsglobal.com/who-we-are for details.

[14] Being *poked* in our setting is thus equivalent to being *treated* in difference-in-differences parlance.

price change. The estimated increase in churn after the notification—but before the changes are implemented—among consumers for whom the change is cost neutral (or cost reducing), is interpreted as evidence of the poking effect. Note that the plans offered by competitors did not change during this period, such that an increased churn among consumers who would be better off is most readily explained by the poke *in itself* having an effect on consumer behavior.

2.3 Impact of the price change on consumers' expenditures

Since our identification strategy is based on grouping poked consumers by the extent to which the price changes affect their costs, we face an obvious measurement problem for consumers who churn or switch plans prior to the actual price change, as we do not observe their realized costs on the plan assigned by the operator. Following Hortaçsu et al. (2017) and Ketcham et al. (2015), we address this problem by calculating the cost given pre-consumption and new prices to obtain hypothetical costs. This gives us an estimate of what the consumer would have paid in the period prior to the poke if subject to the new price structure (i.e., keeping the level and composition of consumption fixed). Next, we construct the variable *Save*, which is the difference between the (monthly average) actual and hypothetical costs during the 6 months prior to notification.[15] Consider the following illustration as an example: if a consumer, on average, paid 20 EUR per month in the pre-period, and that same consumption would cost 10 EUR per month with the new price plan, *Save* equals 10. We consider a consumer as better off if the new price schedule would have lowered costs ($Save > 0$) and worse off if the price change would have increased costs ($Save < 0$).

Note that *Save* is based on each consumer's average consumption during the 6 months prior to the poke (SMS notification), such that its validity depends on pre-consumption being a reasonable approximation to the consumer's counterfactual consumption under the price structure of the assigned plan.[16] Since marginal prices were generally different between the removed and assigned plans, cf. Table A.1 and Table A.2, we would expect that this measure of the extent to which a consumer is better or worse off to be a conservative one, since the consumer should be able to do weakly better (in utility terms) from reoptimizing their usage under the new marginal prices. It is also worth noting that the consumption patterns of the consumers reassigned to new price plans are stable both before and after the price change, indicating that these consumers do not adjust their consumption in response to changes in marginal prices (cf. Figure 1).

[15] $Save_i = average[cost(X_i \mid plan^i)] - average[cost(X_i \mid plan^{new})]$, where $plan^i$ is chosen plan and X_i is average consumption, both of consumer i and prior to the notification. $plan^{new}$ is consumer i's plan after the potential reassignment.

[16] Using past consumption to compare different price plans is analog to "potential savings from switching" in Hortaçsu et al. (2017), "above minimum spending" in Ketcham et al. (2015), and the definition of "maximum potential savings" in Miravete (2003), though the latter use consumption under a flat rate price scheme as measure for maximum consumption.

Figure 1: Average consumption among poked consumers

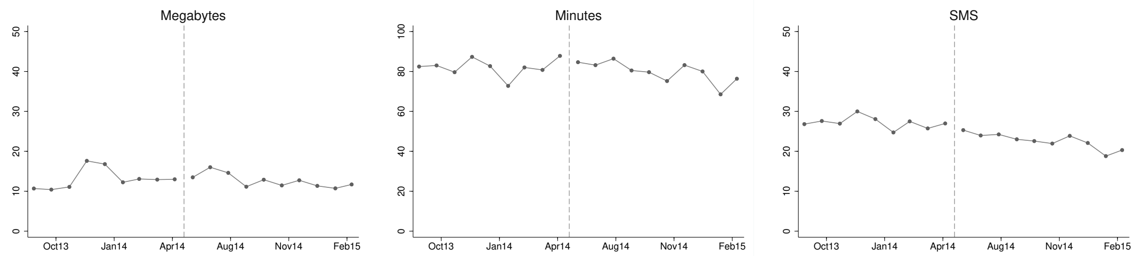

Note: The figures shows the average monthly usage of Megabytes (left panel), calling minutes (mid panel) and SMS (right panel) for consumers who are poked (notified about the price change) and reassigned to new plans. The vertical dashed line indicates the moving date, i.e., the date the consumers face new prices.

2.4 Descriptive statistics

Table 1 provides summary statistics on the consumers who were (potentially) affected by the price change (*Poked*) and those who were not (*Other*). The *Poked* group include consumers who hold one of the reassigned plans at time of poking, and consumers who left the company while holding one of the reassigned plans prior to the time of poking. Residency balances well across poked and other consumers.[17] While, on average, the two consumer groups are different from each other in terms of demographics and mobile usage. This is to be expected, since consumers who have their subscription reassigned have stayed with their price plan for a longer period prior to the reassignment. Note that *tenure* in this table is measured as the number of months a consumer has held their current plan. This explains the notable reduction in tenure for the poked group in the post-period, due to being assigned new plans.

Figure 2 shows the monthly churn rates for poked consumers and consumers on plans that were not reassigned.[18] Before the SMS notification, the churn rate among the poked group is relatively stable and lower than that for other consumers. Notice that although the levels differ, both groups seem to face very similar monthly variations, indicating that they are influenced in the same way by the competitive environment. At the month of the SMS notification, however, churn rate of consumers on notified plans more than triples compared with the month before, staying high for the 3 months from notification until reassignment. There is a noticeable increase in churn rate for the poked consumers in the month prior to notification due to the timing of churn registration in our data.[19] In our following estimates, this tends to overstate churn rate

[17]Place of residency is categorized using the operator's categories.

[18]More precisely, consumers on the reassigned plans. Consumers that churn before the poking are not poked.

[19]Unfortunately, we only have the exact date the consumers left the company for about half of the consumers actually leaving. If we plot churn rate based on only consumers with registered churn dates, no increase in churn rate is observed among the poked consumers before the poking date (cf. Figure A.1). Because of the lack of exact churn dates for the rest of the consumers, a churn is registered as the last month a consumer uses the subscription to call, send an SMS, or connect to the Internet. If a consumer leaves in March before using the phone at all that month, the churn is registered in February. Poked consumers have, on average, lower traffic volumes than non-poked consumers (cf. Table 1). Hence, they are more likely to not have used their phone at all before the *poke*, in the month they were poked.

Table 1: Summary statistics

		Poked		Other	
		Pre	Post	Pre	Post
Residency in:	4 largest cities	0.16	0.16	0.18	0.18
		(0.37)	(0.36)	(0.39)	(0.38)
	City other	0.08	0.08	0.09	0.09
		(0.27)	(0.27)	(0.29)	(0.28)
	Urban other	0.04	0.03	0.01	0.01
		(0.19)	(0.16)	(0.11)	(0.09)
	Village	0.47	0.49	0.45	0.46
		(0.50)	(0.50)	(0.50)	(0.50)
	Other	0.25	0.25	0.26	0.26
		(0.43)	(0.43)	(0.44)	(0.44)
Age		61.78	62.78	47.92	48.51
		(14.52)	(14.07)	(14.25)	(14.19)
Female, share		0.41	0.42	0.50	0.51
		(0.49)	(0.49)	(0.50)	(0.50)
Churn, percent		1.08	1.61	1.22	1.04
		(10.32)	(12.60)	(10.96)	(10.15)
Plan change, percent		0.39	8.59	4.46	4.61
		(6.23)	(28.03)	(20.64)	(20.97)
Tenure, months		49.28	13.02	12.30	13.97
		(33.62)	(24.03)	(12.92)	(12.37)
Save		-6.44	-6.00	-10.07	-10.13
		(9.64)	(9.73)	(19.36)	(19.23)
Fixed fees		1.41	13.19	22.66	25.12
		(2.70)	(6.75)	(9.61)	(8.63)
Variable fees		8.99	8.16	18.48	15.80
		(11.86)	(12.54)	(18.94)	(19.16)
Voice minutes		83.13	99.37	219.03	230.80
		(117.02)	(141.14)	(258.97)	(287.45)
SMS		28.63	30.81	112.17	115.29
		(57.88)	(65.05)	(146.68)	(166.68)
Megabytes		14.83	56.15	443.08	821.32
		(187.57)	(387.27)	(979.31)	(1615.39)
No. of customers		284,627	266,830	665,207	618,669

Note: The poked group are customers who hold one of the reassigned plans at time of poking, or customers who left the company while holding one of the reassigned plans prior to the time of poking. The other are customers on plans which are not reassigned. Pre indicates the six month prior to the poking, while post is the 17 month from the poking month. Standard deviations in parentheses.

in the poked group in the period before notification, thereby giving a downward bias in the estimated increase in churns.

We perform a robustness check to assess the potential for our results being affected by large differences between poked and non-poked consumers, particularly in their tenure.[20] We utilize an event in April 2011, where the operator simultaneously ceased sales of 2 migrated plans and introduced 4 non-migrated plans (see Figure A.3). Development in churn is virtually unchanged if we restrict our poked group to consumers who signed up for the 2 migrated plans in the 6-month period prior to April 2011, and our control group to consumers who signed up for one of the 4 non-migrated plans in the following 6-month period (see Figure A.4).[21] A noticeable difference is the lower overall churn among non-poked customers, making the churn rate of the two groups essentially identical in the pre-period, suggesting that selection involved in longer tenure is the reason for different levels of churn between the groups in the pre-period when using the full sample.

Figure 2: Churn rates (percent)

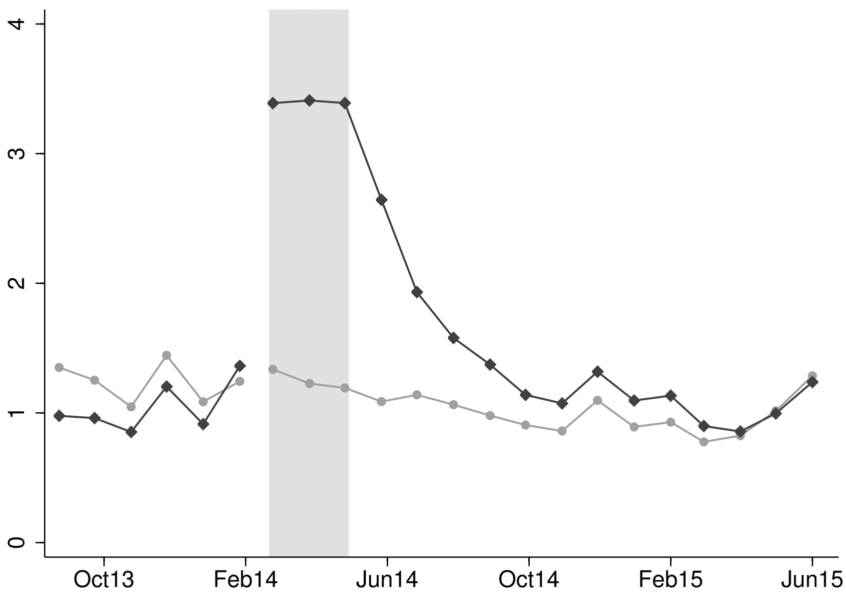

Note: The figure shows the monthly churn rate in percent separately for poked consumers (black line) and non-poked consumers (gray line). The time from poking consumers to their actually being reassigned to a new price plan is indicated by the gray area.

2.5 Market and competitive environment

The telecom markets in Europe are generally characterized by having two to four vertically integrated network operators and numerous service providers. The markets are national, but

[20] See Israel (2005) for a study of the importance and effect of tenure on consumer behavior in the setting of auto insurance.

[21] This reduces our sample from 937,223 individuals (272,222 in the poked group) to 67,025 (17,436 in the poked group), measured in the month prior to notification (February 2014).

subject to increasing regulatory pressure from the European Commission in the direction of reducing market boundaries. Market penetration is generally very high, often measured as exceeding 100 percent (meaning that there are more mobile phones than people). Firms offer both prepaid and postpaid subscription, but the majority of European mobile users are postpaid users, which implies that they subscribe to their telecom provider's service and pay bills on a regular basis.[22] With regard to product quality, service providers are differentiated along several dimensions, where geographic coverage and mobile Internet speeds are the most important. Other dimensions of differentiation include, for instance, customer support quality, mobile phone offers, and bundled digital services (cloud storage, virus protection, music-streaming services, etc.).

Since our identification strategy is based on utilizing a sudden price change, it is important to understand the competitive dynamic, that is, how and whether competitors might respond. In this section, we explain how firms compete for consumers and argue that competitor response is not likely to be of concern.

An important feature of the market dynamics of telecommunications is the existence of both list prices and specialized offers given to consumers. Prices and advertising that are public are often referred to as above-the-line offers, whereas offers given in private to consumers as part of special campaigns are referred to as below-the-line offers. The extent of below-the-line offers in a market typically depends on the intensity of competition and national regulation. Below-the-line offers can be given to both existing consumers and consumers of other suppliers. A supplier has much more information about its current consumers than those of its competitors. This difference in level of information makes the nature of below-the-line offers very different between these two groups: offers to existing consumers can be tailored and specialized, but this is not possible for consumers supplied by competing providers. The market from which we have data has a relatively low extent of below-the-line offers given to consumers from competing suppliers.

This dynamic in above- and below-the-line pricing is important for how competitors might respond to the price change studied here. Competitor response is a concern since it could drive consumer churn and thereby weaken identification. We believe that competitor response is not likely to be of concern. First, with regard to above-the-line prices, the relevant competitors do not change their prices during the relevant period of our study. Second, with regard to below-the-line prices, the consumers affected by the price change constitute a select group of consumers who are not identifiable by competing providers wanting to approach them with offers. We have also discussed this price change with the responsible managers to ask whether they were aware of any below-the-line campaigns from competitors in response to the price change. Managers typically have very detailed knowledge about what competitors are doing through real-time consumer feedback. The responsible managers were not aware of any campaigns started as a direct response to the price change.

[22]Prepaid services usually do not have regular payments, and usage is limited by the extent to which the consumer has paid in advance for megabytes, SMS, and voice calls.

3 Estimating the poking effect

It is conceivable that many consumer decisions are based on habits, due to, for example, time and attention constraints. Past purchase behavior may therefore be repeated without much consideration. Such a pattern of repeating past behavior can extend until something special happens or catches the consumer's attention—what we refer to as a poke—which forces them to make an active choice, potentially reconsidering their previous decision. Put differently, their choice of service provider is closed unless given a reason to reconsider, i.e., reopen the decision, and make an overall assessment of the offers available on the market. There are several kinds of events that could poke a consumer to reopen their choice of provider, such as advertising campaigns, an unexpectedly high bill, changes in prices, or some other event making their subscription to the service more salient. In this paper, we focus on estimating the poking effect from sending a notification of the price change by SMS.

3.1 Difference-in-difference estimation

From Figure 2, we see that there is a strong increase in churning following notification of the contract change. This should not come as a surprise since a large proportion of the consumers would experience an increase in monthly costs with the new prices. To elicit the poking effect caused by the notification, we are particularly interested in the response of consumers whose expected cost would decrease or remain unchanged. To this end, we estimate a difference-in-difference model where we split consumers into discrete groups of predicted savings (*Save*), and estimate how the notification affects the churn rate within groups of consumers who will be relatively better or worse off after the change in terms of cost.

We define four time periods: *before*, *during*, and two *after* periods. *Before* contains the 6 months prior to the notification about the price change for the treated consumers. *During* covers the months March to May 2014, where March is the month of notification, and May is the month where consumers are reassigned. The *after* periods split the remaining time into two 8-month windows. Our goal is to investigate the extent to which the change in expected expenditures can explain the increase in churn following the price change, with particular focus on the *during* period. For ease of exposition, we ignore the *after* periods in the following presentation; however, they are included in the estimation.

We then split consumers into 12 categories depending on their change in expenditures (i.e., the value of *Save*).[23] For consumers who were not poked, *Save* measures the savings they would have obtained *if* had they been reassigned or chosen the new price plan themselves. For each category of *Save*, we estimate the following model:

$$churn = \gamma d_{poked} + \lambda d_{during} + \delta d_{during} d_{poked} + \boldsymbol{x}'\boldsymbol{\beta} + \epsilon, \quad (1)$$

where d_{during} is an indicator for the during period, i.e., the months following the SMS notifica-

[23] In practice, our model could be considered and formulated as a triple difference specification, with a difference in DiD changes between groups of consumers with different levels of predicted savings.

tion; d_{poked} is an indicator equal to one if the consumer received a SMS notification; and $\boldsymbol{x}$ is a vector of controls, containing age, gender, and area of residence.[24]

We are particularly interested in δ, as it measures the change in churn for the *poked* group in the months following the SMS notification compared with before the price change. For our estimates to be valid, we need the assumption of common trends to hold within each savings group. In our sample, pre-treatment trends are very similar for each group (see Figure A.2 in the Appendix), thus making us conclude that the control group provides us with a reasonable approximation to what the churn rate in the poked group would have been in the absence of the notification.

3.2 Empirical results and discussion

Figure 3 shows the estimated increase in churn, δ from Equation 1, i.e., the increase in churn in the *during* period compared with the *before* period for the *poked* consumers per category of predicted *Save*. Not surprisingly, we see that the increase in churn is larger with lower *Save*. However, the increase in churn is also substantial for consumers predicted to be unaffected or better off under the new price plan. Among the consumers predicted to save from the new price plans, churn increases by 0.5 to 1 percentage points, which amounts to an increase of 50 to 80 percent relative to the *before* churn rate.[25]

Furthermore, we do not find evidence of heterogeneity in the poking effect within areas stated as important drivers of behavior in the previous literature; salience of specific prices (cf. Grubb, 2015a; Bordalo et al., 2012; Gabaix and Laibson, 2006),[26] or salience of the size of a bill (cf. Grubb, 2015b).[27] On the other hand, we do find substantial heterogeneity in the poking effect if we classify the consumers by age, cf. Figure A.5. Overall, younger consumers are more prone to churn and respond to the price change.

We denote the increase in churn as a poking effect, though we stress that this is merely a description of behavior. For managers, knowing about this effect is valuable, as it helps them

[24]More precisely, the formulation of our model is

$$churn_{it} = \gamma d_{poked} + \sum_{t \in \mathcal{T}} \left(\lambda_t d_t + \delta_t d_t d_{poked} \right) + \boldsymbol{x}_{it}' \boldsymbol{\beta} + \epsilon_{it},$$

where $\mathcal{T} = \{during, after_1, after_2\}$. In principle, we could estimate time and treatment effects month by month. This does not change our conclusions in any meaningful way (results are available from the authors upon request). For ease of exposition, we divide time into broad groups, and report results on the *during* period exclusively.

[25]We cluster standard errors by the price plan the consumer holds prior to the poke and the month in which the consumer entered this price plan. Thus, consumers entering the same price plan in the same month are in the same cluster, i.e., consumers who entered a price plan under similar market conditions. If we were to cluster observations by consumer instead, all standard errors would be smaller.

[26]The change in the monthly fixed fee can be expected to be more salient than the other price changes. Hence, we could assume consumers who experience a larger increase in the fixed fee to be more prone to churn. However, we do not find any difference in the poking effect for consumers facing a large compared with a small increase in the fixed fee.

[27]All consumers experience an increase in the fixed fee, cf. Table A.1 and Table A.2. In the telecom industry, bill shocks—an unusually high bill—are perceived to increase the propensity to churn (c.f. Grubb, 2015b). Even so, in the *before* period, bill shock does not correlate with churn, and we do not find evidence for a recent bill shock to explain the churn increase in the *during* period among the *poked* consumers.

Figure 3: Estimated increase in churn rates (in percent) for the *poked* in the *during* period by predicted savings categories

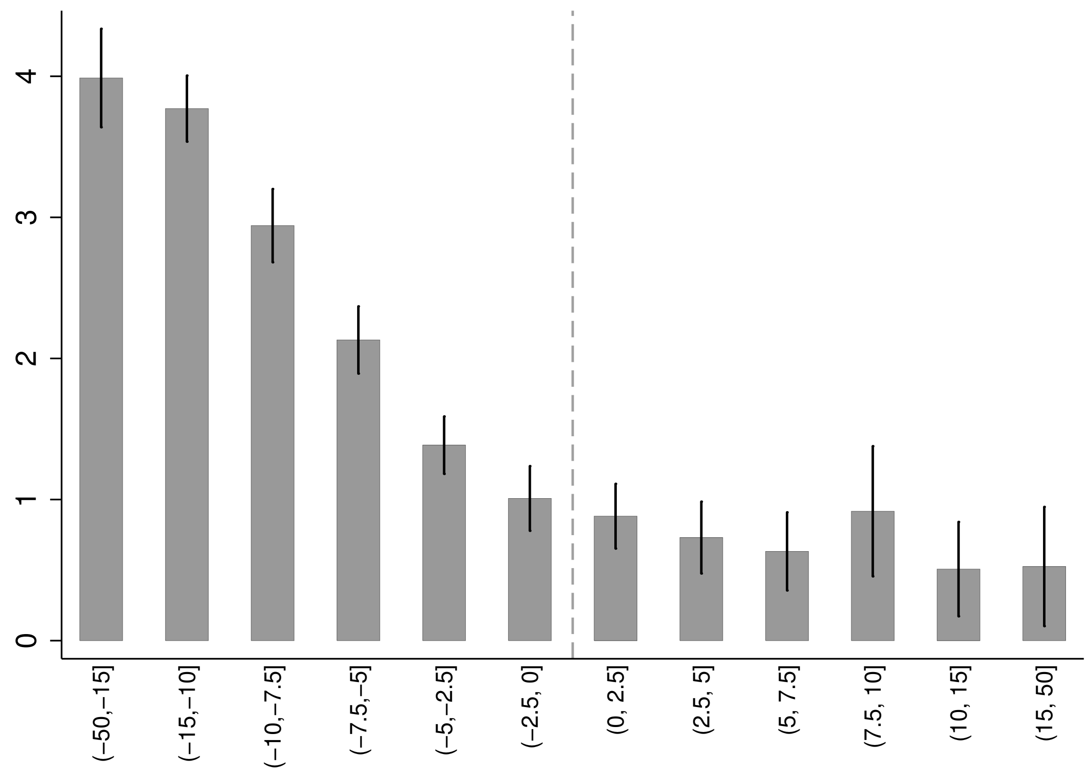

Note: The figure shows estimated churn increases, in percentage points, from the before to during period (δ) for different intervals of predicted savings categories with 95 percent confidence intervals indicated by error bars. Robust standard errors clustered at the individual's choice of price plan before the poke, i.e., a combination of price plan and the tenure on that price plan (3,089 clusters).

understand the potentially unintended consequences of making changes in contract terms for existing customers.

3.3 Survey evidence and discussion of the role of inattention

Having established the existence of a poking effect, what can we infer about the underlying mechanisms? Given the nature of the intervention—an unexpected SMS that informs consumers that prices will change—it seems natural to assume that the effect is due to some form of limited attention. Consumers have limited time available, and it is time consuming to constantly keep track of alternative offers in the market. The SMS could serve as a signal (or poke) that catches their attention and suggests that it is time to reopen their choice of telecom provider and get an overview of alternative offers available.

However, there are other explanations which also are consistent with the poking effect: Since service providers have an incentive to increase profits at the expense of consumers (when possible), it is rational for consumers to view price changes with skepticism. Hence, some consumers—even when they would be better off—may automatically assume that a price change leads to higher costs, and therefore switch providers. Alternatively, some consumers may miscalculate the effect of the price change on their costs and switch due to an erroneous belief that costs will increase. In this section we look closer at the importance of these explanations.[28]

With regards to miscalculating the effect of the price change on costs, there are two pieces of evidence which point in the direction that this is not an important factor. First, among the poked consumers who churned but whose costs would have decreased with the new prices, we find that about 80 percent leave for an operator offering a plan that will reduce their costs even further. This can be seen in Figure 4.[29,30] The figure hence indicates that most consumers are neither uninformed nor unable to make cost-minimizing decisions in this market. Hence, it is likely that most of these consumers are able to understand that the new prices would lower their cost. Second, the poked consumers generally has a lower level of mobile phone usage than the average consumer (cf. rows "Voice minutes", "SMS" and "Megabytes" in Table 1), implying that the complexity of calculating the impact of the price change is low.

But, as previously mentioned, consumer may automatically assume that the price change is bad news, no matter how good it may appear. This heuristic is arguably rational, as it holds true for more than 50 percent of the poked consumers. In order to investigate the importance of this explanation relative to the role of limited attention, we conducted a survey on a sample of 451 customers of the telecom company.

The first question in the survey use the frequency of reoptimization as a measure of attention. We asked *How often do you check the prices of other mobile service providers than your current provider?* About 60 percent of respondents check prices rarer than every second year, indicating that most consumers are fairly inactive with regards to keeping track of prices. Hence, inattention is likely to be an important source of consumer inertia in this market. We calculate that around 7% of the consumers check prices in any given month, based on the answers in this survey.[31]

The second question assesses how an announced price decrease would be interpreted by consumers. We asked *Imagine that your current mobile service provider lowered the price on your price plan. What best describes your reaction?* We had four different answers that the respondents could choose from based on two dimensions: trust/distrust and poke/no poke. The responses

---

[28] There are other mechanisms that could potentially interact with our main effect in important ways, such as learning from experience (see, e.g. Miravete and Palacios-Huerta, 2014b). In our setting, this would imply that consumers decide to churn based on learning how the new prices affect them. Since the poking effect we identify is at its strongest immediately after notification—*before* the consumer has had the time to experience new prices—learning cannot be the driver of the most dramatic changes in behavior we document here. It is possible that learning plays some role in the elevated churn we see for poked consumers in Figure 2, in the period *after* the change is implemented (from May 2014 and onwards), though this effect is much less pronounced than the increase prior to the change. Hence, the timing of the effects rules out learning as the driver of the poking effect.

[29] The figure only includes consumers who churn in the 3-month period after being informed about the price change ($N = 10,839$); in other words, prior to the price change being effectuated.

[30] At the time of the poking, competitors offer price plans that dominate those held by the poked consumers, cf. Table A.3 versus Table A.1 and Table A.2.

[31] We assigned a (monthly) probability of checking prices to each respondent equal to the inverse of the reported duration (in months), and setting the duration of "Less often" (than every second year) to every fourth year. This means that those who report to check prices monthly are assigned a probability of 1, while it is 1/3 for those who report to check quarterly, etc.

Figure 4: Cost difference between new plan and competitor's lowest-cost plan

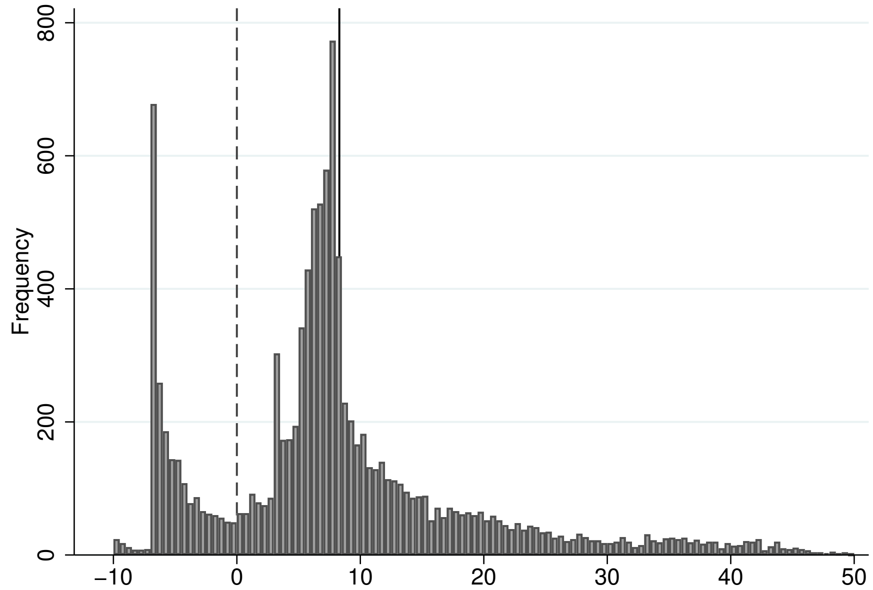

Note: The figure shows, limiting the sample to consumers who will experience reduced expenditures with the new prices and churn in the during period ($N = 10,839$), the histogram over each consumer's difference between the cost of their new assigned plan compared with the lowest expenditure plan offered by the operator they churn to. The solid vertical line indicates how much these consumers, on average, save by churning if they chose the competitor's least-expensive plan. The dashed vertical line at zero separate consumers who leave for operators that do not offer less expensive plans (left side) from consumers who leave for operators offering less expensive plans (right side).

were as follows (our classifications added in bold).

- **Trust/no poke:** *I would be glad and not thought more about it. I want to spend as little time as possible on prices*
- **Distrust/no poke:** *I would find it suspicious, but not thought more about it.*
- **Distrust/poke:** *I would find it suspicious and used the opportunity to check the prices of competing providers*
- **Trust/poke:** *I would be glad, but used the opportunity to check the prices of competing providers*
- *Don't know*

The results are shown in Figure 6. It it immediately clear that distrust is rare when facing a price decrease, and that a significant share of consumers are poked: 19 percent would welcome the change, but at the same time used the opportunity to check competing offers. Together with the answers to the frequency of price search (first question), we calculate that roughly 25% of customers will check prices following such a poke.[32]

[32]Calculated by assigning a probability of 1 to check prices for consumers reporting that they would use the opportunity to check the prices of competing providers, while assigning a probability to check prices equal to the inverse of the reported frequency of search from the first question (see Footnote 31 for details).

Figure 5: Frequency of competitor price comparison

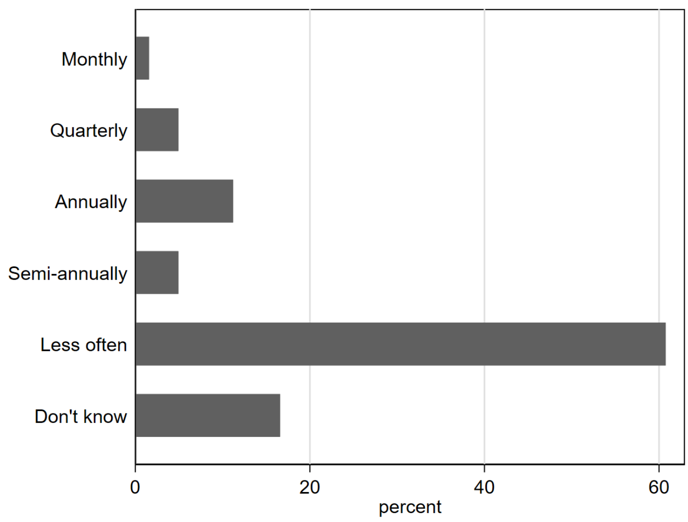

Note: The figure shows survey results of the question "How often do you check the prices of other mobile service providers than your current provider?" 446 respondents, in which all are customers of the relevant telecom operator, answered the question. The respondents chose between the categories shown in the figure. Five respondents, who did not answer, are not included in the figure.

Lastly, we wanted to explicitly measure the consumers' level of trust/distrust in their mobile service provider. The respondents were asked *To what extent do you agree with the following statement: "I can trust that my mobile service provider does not try to trick me with hidden prices and fees"*. Only 18 percent disagree with this statement, which we interpret as evidence that distrust is not widespread among consumers.

Taken together the survey responses provide indicative evidence that i) inattention is common among consumers in this market, ii) a price decrease would poke a significant share of consumers to reoptimize, and iii) distrust is not widespread. These responses point to attention playing an important role in explaining consumer inertia in this market. Moreover, the poking effect is most likely driven by an increase in attention.

Figure 6: Consumer response to a price decrease

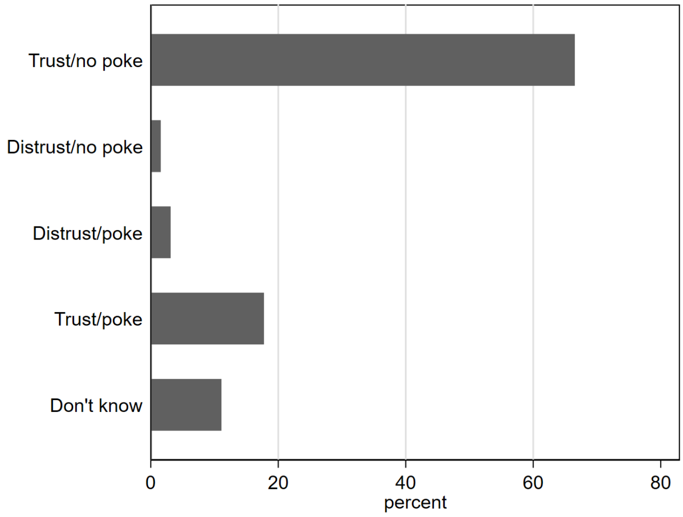

Note: The figure shows survey results of the question "Imagine that your current mobile service provider lowered the price on your price plan. What best describes your reaction?" 447 respondents, in which all are customers of the relevant telecom operator, answered the question. The respondents chose between the categories shown in the figure. Four respondents, who did not answer, are not included in the figure.

4 A model of consumer attention and choice

In this section we estimate a model of consumer attention and choice. The main intuition behind the model is based on the notion that a consumer can be either attentive or inattentive. An inattentive consumer is on "autopilot", i.e., automatically repeating past behavior without further thought. An attentive consumer is involved, in the sense of actively optimizing by considering the alternatives on the market. Thus, a churn event is a combination of becoming attentive (i.e., considering alternative offers) and choosing a plan from a different provider. At the same time, there will be consumers who are attentive but decide not to churn since their best option is with their current company (possibly also staying on their current plan). An important feature of our model is that it allows us to separate the effect of the poke on consumer attention from the effect of removing consumers' current plans.

The model can be described as a two-stage process of consumer choice. In Stage 1, the consumer either becomes attentive or stays inattentive. In Stage 2, conditional on being attentive, the consumer chooses between the plans currently offered on the market. Thus, churn happens when the consumer is attentive and prefers one of the plans offered by competing providers over those offered by the current provider. Potential "implicit brand effects" will have an influence on decisions in Stage 2: Consumers who are attentive but still do not churn in the presence of less expensive alternatives on the market are willing to pay a brand premium.

When consumers are informed about a price change, their propensity to churn increases, already 2 months before the price change is effectuated. This behavior is consistent with consumers behaving forward-looking in the following way: Imagine making a choice knowing

Figure 7: Mobile operator trust and hidden prices

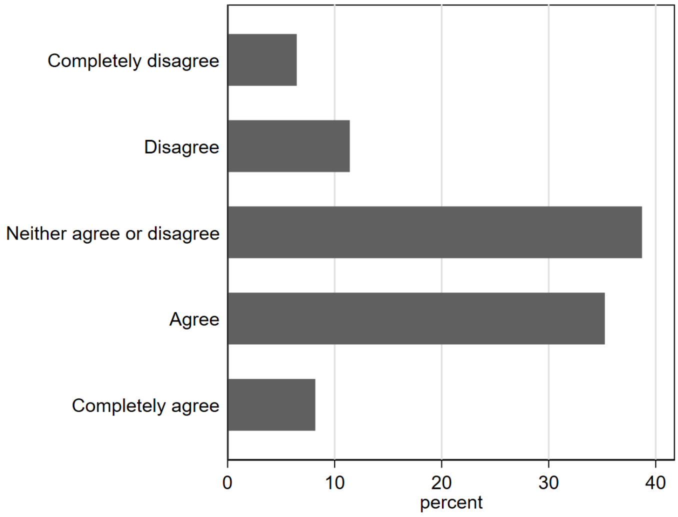

Note: The figure shows survey results of the question "To what extent do you agree with the following statement: "I can trust that my mobile service provider does not try to trick me with hidden prices and fees". 449 respondents, in which all are customers of the relevant telecom operator, answered the question. The respondents chose between the categories shown in the figure. Two respondents, who did not answer, are not included in the figure.

there is only a certain probability that you can revise your choice at any given point in the future, e.g., due to inattention. When you learn that the payoff of a particular choice will change in the future, this should be factored into your current decision. This type of forward-looking behavior implies that consumers increase their propensity to churn when they are informed about an upcoming price change, even though they would have preferred their initial choice of price plan in any given period (e.g., if the plan would not change in the future, if consumers could freely choose at any point in time, or if they were myopic).

4.1 Model specification

Formally, we estimate a model of the joint probability of being attentive and the optimal choice of plan conditional on being attentive. Let p^A_{it} be the probability that consumer i is attentive in month t, and let $p^C_{it}(j)$ be the probability that consumer i chooses plan $j \in \mathcal{J}_t$ in month t (conditional on being attentive), where $\mathcal{J}_t$ is the set of plans offered by all firms in month t.[33] Let y_{it} denote consumer i's choice of plan in month t. The probability of observing a choice $y_{it} \neq y_{i,t-1}$ is $p^A_{it} p^C_{it}(y_{it})$, whereas the probability of observing $y_{it} = y_{i,t-1}$ is $p^A_{it} p^C_{it}(y_{i,t-1}) + (1 - p^A_{it})$. In the case where a consumer changes plans from one period to the next, the consumer would necessarily be attentive, whereas the consumer repeats the previous choice either if i) they are attentive *and* it is the optimal choice (which occurs with probability $p^A_{it} p^C_{it}(y_{i,t-1})$); or ii) they

[33]Note that we do not include an outside option in the choice set, since it is not relevant in our case. The market penetration of mobile services is nearly full in the market considered here, and we do not find any evidence that consumers stop using mobile services when they churn. We therefore assume that consumers choose one of the plans offered on the market in our model.

are inattentive (which occurs with probability $1 - p_{it}^A$). This basic structure of the model is equivalent to the one utilized by Hortaçsu et al. (2017) to study consumer inertia in the Texas retail electricity market. Similarly, we take a reduced-form approach to modeling p^A, specifying it as a logit depending on age, gender, and centrality of area of residence of the consumer. However, our setting provides additional, identifying variation through the addition of time dummies interacted with the treatment status of the consumer, which captures the effect of the poke as a shifter of attention. We specify the probability of becoming attentive as a standard binary logit:

$$p_{it}^A = \frac{e^{f_A(W_i, d_{poked}, t)}}{1 + e^{f_A(W_i, d_{poked}, t)}}$$
$$f_A(W_i, d_{poked}, t) = W_i'\eta + \gamma d_{poked} + \sum_{\tau=1}^{T} \lambda_\tau d_\tau + \sum_{\tau=1}^{T} \delta_\tau d_\tau \cdot d_{poked}, \tag{2}$$

where W_i is the vector of individual characteristics listed above; d_τ is an indicator for t being the τ^{th} of the sample; and d_{poked} is an indicator for consumers who, at any point prior to the notification (in March 2014), are on a plan that will be moved, or who were on such a plan when the notification was sent (regardless of whether they switch plans later). The notification sent out by the firm thus provides a useful source of variation. When we condition on the (potential) change in payoff from staying with the current choice at the time when the consumer learns of the upcoming change, the notification itself allows us to quantify consumers' attention separately from their choice behavior.

We let the flow utility of consumer i from plan j in month t be given by

$$v_{ijt} = \alpha c_{ijt} + \sum_f \kappa_f d_{jf}, \tag{3}$$

where c_{ijt} is the cost of the plan to the consumer and d_{jf} are indicators for plan j being sold by firm f. The pattern we find in our reduced-form analysis is consistent with consumers being inattentive and forward-looking, and we therefore chose to model them as accounting for future inattention when being attentive and making an active choice. A consumer will expect the duration until the next time they make an active choice to be the reciprocal of the attention probability $1/p_{it}^A$, and the relevant value to the consumer when deciding between the plans, disregarding discounting, is $V_{ijt} = v_{ijt}/p_{it}^A$.[34] Under circumstances where no changes in the contract structure of any plans are expected, the maximum of the current period utility (given by v_{ijt}) and that of the expected flow utility until the next active choice (given by V_{ijt}) coincide. This distinction is relevant to consumers who are informed of the upcoming change in contract *and* who become attentive before that change is enacted.[35] Even if their current plan is optimal

[34]Note that this expression is only valid if p_{it}^A (and v_{ijt}) is fixed over time, or if we assume that the consumer expects their current "level of attention" to last into the foreseeable future. Otherwise, the correct expected duration is $\sum_{\tau=1}^{\infty} \tau p_{i,t+\tau}^A \prod_{s=1}^{\tau-1} \left(1 - p_{i,t+s}^A\right)$. Since the differences in terms of implications for choice behavior in our model should be small, we opt to use the simplified term in our analysis. Furthermore, it would not be clear how we should extend p_{it}^A out of the time frame of our sample, in addition to the assumption of consumers having perfect foresight of *changes* in future inattention not necessarily being appealing.

[35]See Appendix B for a derivation of the relevant choice values from dynamic optimization of the consumer's

given the price structure of plans offered in the market at that time, these consumers might decide to choose a plan that is worse than their current one to avoid being stuck with the new conditions because of inattention.

Since the notification of the change comes 2 months prior the change itself, consumers who become attentive at this point will get at most 2 months on the current plan before their conditions change. Letting t_C denote the month of the change (March 2014, where treated consumers will be reassigned to the new plan), the value of choosing the current plan 2 months before the change is

$$V_{i,y_{i,t_C-3},t_C-2} = \left[1 + (1 - p^A_{i,t_C-2})\right] v_{i,y_{i,t_C-3},t_C-2} + \frac{(1 - p^A_{i,t_C-2})^2}{p^A_{i,t_C-2}} v_{i,y_{i,t_C-3},t_C},$$

where y_{i,t_C-3} denotes the current choice (from the perspective of making a choice 2 months before the change); v_{i,y_{i,t_C-3},t_C-2} is the current (flow) payoff on the current plan; and v_{i,y_{i,t_C-3},t_C} is the payoff under the contract terms after the change. The first term is the expected value of one "voluntary" month on the current plan and an additional month on this plan because of inattention. The second term is the expected value from not making an active choice until the new contract terms take effect and being stuck with the new prices because of inattention. Similarly, for a consumer who is notified and is attentive *one* month before the change, the value of choosing the current plan is

$$V_{i,y_{i,t_C-2},t_C-1} = v_{i,y_{i,t_C-2},t_C-1} + \frac{1 - p^A_{i,t_C-1}}{p^A_{i,t_C-1}} v_{i,y_{i,t_C-2},t_C},$$

where there is now 1 month actively chosen on the current plan, before inattention might lock the consumer in with the new terms with a probability $1 - p^A_{i,t_C-1}$ in each period.

We model the full (expected) utility of choosing plan j as

$$u_{ijt} = V_{ijt} + \epsilon_{ijt}$$

where ϵ_{ijt} is an i.i.d. Type I Extreme Value random utility term. Thus, the choice probabilities (conditional on attention) are given by

$$p^C_{it}(j) = \frac{e^{V_{ijt}}}{\sum_{k \in \mathcal{J}_t} e^{V_{ikt}}} \tag{4}$$

We assume that all consumers actually choose a mobile subscription, i.e., that no consumers stop having a mobile subscription.[36]

The likelihood of observing consumer i's choice y_{it} in period t is then

$$L_{it} = p^A_{it} p^C_{it}(y_{it}) + (1 - p^A_{it})\, \mathbb{1}(y_{i,t-1} = y_{it}),$$

problem.

[36]For the discrete choice model, this implies that we do not have a separate outside option.

where the last part corresponds to the probability of observing a repeated choice due to inattention ($\mathbb{1}(y_{i,t-1} = y_{it})$ is an indicator for the observing the same plan in t and $t-1$ for consumer i).

The log-likelihood of the sample is given by

$$\ln \mathcal{L} = \sum_i \sum_t \ln L_{it},$$

which we maximize with respect to the model parameters.

After estimating the parameters of the model, we can calculate the probability of churn as follows: Let $\mathcal{D}_t$ denote the set of plans offered by competing providers in month t. Choosing a plan from this set means that the consumer churns. The probability of churning is thus given by

$$\Pr_{it}(churn) = p_{it}^{A} \cdot \sum_{j \in \mathcal{D}_t} p_{it}^{C}(j), \tag{5}$$

which we can use to assess the model against what we observe in the data. Note that our choice model does not directly target churn; the probability of churn conditional on attention ($\sum_{j \in \mathcal{D}_t} p_{it}^{C}(j)$) depends on how well the model approximates consumer choice over the full set of contracts.

4.2 Results from model estimation

Figure 8 presents the main results from Stage 1.[37] There are several interesting things to note from the figure. First, the estimated share of consumers who are attentive in a month is fairly low, around 5 percent. This is relatively close to the 7% share of attentive consumers implied by the survey responses (see Section 3.3). Second, the text message informing about an upcoming price change (the poke) had a strong effect on consumers' attention. The share of consumers becoming attentive more than doubled following the poke. This is in line with the strong response to the poke implied by the survey responses, though much smaller. We find it encouraging that our model estimates are fairly in line with the survey evidence, especially since we find it likely that the (implied) attention rate and poking effect are somewhat exaggerated in the survey due to potentially rosy beliefs about own behavior (particularly for reported increases in price search due to a hypothetical price decrease).

Note also that attention varies strongly with age, the base attention rate being nine percent for consumers aged between 20 and 30 years and four percent for consumers aged more than 80 years (see Figure 9). We also find statistically significant, though relatively modest, coefficients the use of calls, text, and data. The results suggest that consumers who on average use more calling and texts are less attentive, whereas the opposite is true for data use. Since the variables are standardized to unit variance, the estimated coefficients imply that the impact (in log-odds) of a one standard deviation increase in usage is equivalent to between one-half and one-third of the impact of going from 20 to 30 years of age.

[37] See the Appendix, Table A.6 for a table of coefficients.

Figure 8: Predicted effect of poke on attention

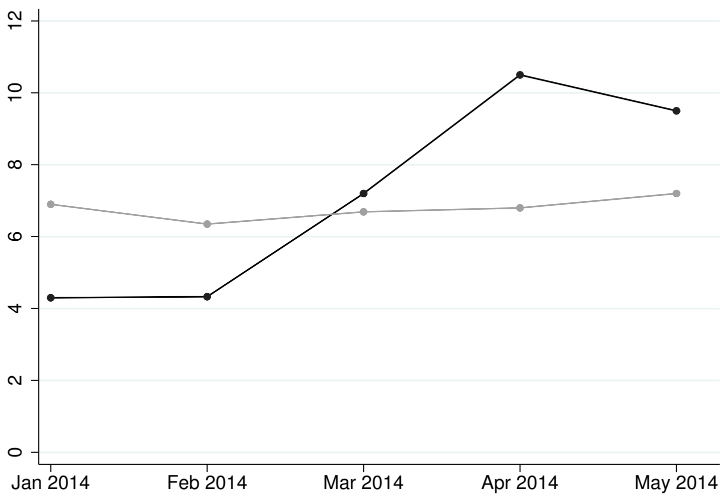

Note: The figure shows the estimated levels of attention from the model, in percent of consumers being attentive. The black line indicates the poked consumers, while the gray line indicates those not poked. The model allows attention to vary separately for poked and non-poked consumers in each month, in addition to demographics and mobile usage patterns.

Figure 9: Attention by age

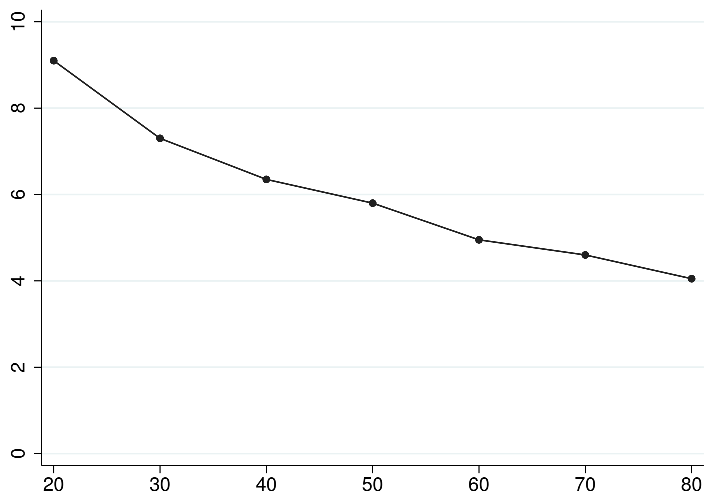

Note: The figure shows the average predicted levels of attention, in percent of consumers, from the estimated model across 10-year age groups (left side of interval provided) for the period prior to the poke.

In Figure 10, we compare the predicted churn from our model (Equation (5)) with the actual churn observed in our data. This exercise serves as a validation that our dynamic model provides a good description of actual behavior. Though the predicted levels of churn are about 0.5 percentage points lower than the actual levels for the poked consumers and about 0.25 percentage points higher than the actual levels for the non-poked consumers, the estimated changes over time closely match the observed patterns.

Figure 10: Predicted effect of poked on churn

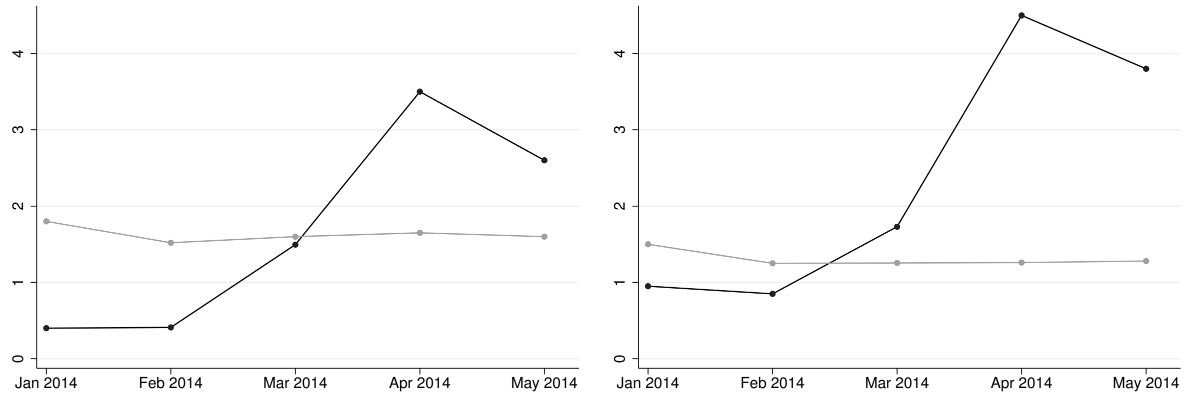

(a) Predicted (b) Observed

Note: The figures show the predicted effect of the poke on churn, in percent, for the poked consumers (black line) and those not poked (gray line). Panel (a) shows the implied rate of churn from the model, while panel (b) shows the observed rate of churn in the subsample used for estimation. The model allows attention to vary separately for poked and non-poked consumers in each month, in addition to demographics and mobile usage patterns.

We also obtain an interesting result on brand valuation from the estimation of Stage 2. This effect can be thought of as the average monetary value of staying with the current provider instead of churning. From the coefficient estimates, we find this brand value to be roughly 10 EUR.[38] Put differently, for many consumers, there are cheaper options available on the market that are not chosen to the extent we should expect from a naive model focusing only on the monetary cost of usage. Hence, there are other aspects of the providers, perhaps related to quality of service or pure brand preference, that influence consumer choice (conditional on being attentive). Furthermore, a model ignoring inattention might severely overestimate brand valuation, due to consumer inertia being attributed to loyalty rather than choice frictions.

5 Concluding remarks

We document a significant and strong increase in consumer churn following a notification about a future price change. This effect also holds for consumers whose costs become lower with the new prices. Taken together, our results point to the existence of a poking effect: the

[38]See the estimated coefficients in Table A.6. Note that the scale of the cost variable is in units of 10 EUR. The brand value is calculated by assuming that the marginal utility of income is equal to (the negative of) the coefficient on cost (see, e.g., Train, 2009, p. 39), while the additional utility of choosing the current provider is equal to the difference between the coefficients on the current provider and other providers, yielding

$$WTP_{brand} = \frac{\beta_{current} - \beta_{other}}{-\beta_{cost}} = \frac{0.84 - 0.22}{0.64} = 0.97,$$

i.e., 9.7 EUR brand value of the current provider compared with other providers.

notification about the price change causes many consumers to "wake up" and start searching for alternative offers. We suggest that the main mechanism behind the poking effect is a reduction in inattention, i.e., that consumers have a low probability of actually considering alternative offers, while the notification increases their attention to the possibility of purchasing a different product. We also conduct a survey which provides supporting evidence on the important role of inattention and how a poke might trigger search for alternative offers.

Using a two-stage discrete choice model we then estimate the extent of consumer attention and how it was affected the poke. We find that, on average, about five percent of consumers are attentive in the months leading up to the price change, and that the poke causes the proportion of consumers becoming attentive to more than double. This result has interesting policy implications, as it shows that potentially simple interventions can have a significant effect on consumer switching through increasing attention.

Our results have managerial implications, in particular for providers in markets where consumer involvement is relatively low, such as telecom, insurance, and banking. In such markets consumers tend to be passive for longer periods, until some special event triggers them to re-optimize. Managers should be aware that changes in terms and conditions, and even a price decrease, could trigger churn as it causes many consumers to become attentive and "reopen" their choice of provider. Our findings also suggest that managers could mitigate this effect by making sure that the new terms and conditions are not only better than before, but competitive compared to the new offers that have entered the market since the last time terms and conditions were updated. Lastly, our results show that inattention is not necessarily static and unchanging, even within individuals, suggesting that it is possible to reduce this potentially important impediment to capturing new consumers through targeted measures.

References

ASCARZA, E., R. IYENGAR, AND M. SCHLEICHER (2016): "The Perils of Proactive Churn Prevention Using Plan Recommendations: Evidence from a Field Experiment," *Journal of Marketing Research*, 53, 46–60.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2012): "Salience in Experimental Tests of the Endowment Effect," *American Economic Review*, 102, 47–52.

——— (2013): "Salience and Consumer Choice," *Journal of Political Economy*, 121, 803–843.

DUBE, J.-P., G. J. HITSCH, AND P. E. ROSSI (2010): "State dependence and alternative explanations for consumer inertia," *The RAND Journal of Economics*, 41, 417–445.

GABAIX, X. AND D. LAIBSON (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, 121, 505–540.

GRUBB, M. D. (2015a): "Behavioral Consumers in Industrial Organization: An Overview," *Review of Industrial Organization*, 47, 247–258.

——— (2015b): "Consumer Inattention and Bill-Shock Regulation," *The Review of Economic Studies*, 82, 219–257.

HANDEL, B. R. (2013): "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 103, 2643–2682.

HANNA, R., S. MULLAINATHAN, AND J. SCHWARTZSTEIN (2014): "Learning Through Noticing: Theory and Evidence from a Field Experiment," *The Quarterly Journal of Economics*, 129, 1311–1353.

HO, K., J. HOGAN, AND F. S. MORTON (2017): "The impact of consumer inattention on insurer pricing in the Medicare Part D program," *RAND Journal of Economics*, 48, 877–905.

HORTAÇSU, A., S. A. MADANIZADEH, AND S. L. PULLER (2017): "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market," *American Economic Journal: Economic Policy*, 9, 192–226.

ISRAEL, M. (2005): "Tenure Dependence in Consumer-Firm Relationships: An Empirical Analysis of Consumer Departures from Automobile Insurance Firms," *The RAND Journal of Economics*, 36, 165–192.

KARLAN, D., M. MCCONNELL, S. MULLAINATHAN, AND J. ZINMAN (2016): "Getting to the Top of Mind: How Reminders Increase Saving," *Management Science*, 62, 3393–3411.

KETCHAM, J. D., C. LUCARELLI, AND C. A. POWERS (2015): "Paying Attention or Paying Too Much in Medicare Part D," *American Economic Review*, 105, 204–33.

MADRIAN, B. C. AND D. F. SHEA (2001): "The power of suggestions: Inertia in 401(k) particiaption and savings behavior," *The Quarterly Journal of Economics*, 116, 1149–1187.

MARZILLI ERICSON, K. M. (2014): "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange," *American Economic Journal: Economic Policy*, 6, 38–64.

MIRAVETE, E. J. (2003): "Choosing the Wrong Calling Plan? Ignorance and Learning," *American Economic Review*, 93, 297–310.

MIRAVETE, E. J. AND I. PALACIOS-HUERTA (2014a): "Consumer Inertia, Choice Dependence, and Learning from Experience in a Repeated Decision Problem," *The Review of Economics and Statistics*, 96, 524–537.

——— (2014b): "Consumer Inertia, Choice Dependence, and Learning from Experience in a Repeated Decision Problem," *The Review of Economics and Statistics*, 96, 524–537.

SAMUELSON, W. AND R. ZECKHAUSER (1988): "Status quo bias in decision making," *Journal of Risk and Uncertainty*, 1, 7–59.

TRAIN, K. (2009): *Discrete Choice Methods with Simulation*, Cambridge University Press.

Appendix A Supplementary analyses

Table A.1: Poked plan overview Migration 1

Plan	Prices				
	Fixed fee	Minutes included	Minute price above incl.	SMS price above incl.	MB price
Old plan A	8.9	60	0.09	0.07	1.25
Old plan B	4.9	20	0.18	0.07	1.25
Old plan C	2.9	0	0.05	0.05	1.25
Old plan D	4.9	0	0.09	0.07	1.25
Old plan E	0	0	0.05	0.05	1.25
Old plan F	0	0	0.05	0.05	1.25
Old plan G	5.9	0	0.10	0.07	1.25
New plan 1	12.9	100	0.05	0.05	1.0

Note: The table gives an overview of prices of seven old plans (old plan A–G) and the new plan (New plan 1) that these seven old plans are moved to. Standardized currency.

Table A.2: Poked plan overview Migration 2

Plan	Prices				
	Fixed fee	SMS included	Minute price above incl.	SMS price above incl.	MB price
Old plan H	2.9	0	0.07	0.04	1.25
Old plan I	0	0	0.09	0.07	1.25
Old plan J	6.9	0	0.16	0.07	1.25
Old plan K	0	0	0.05	0.05	1.25
Old plan L	0	0	0.05	0.05	1.25
Old plan M	0	0	0.04	0.04	0.50
New plan 2	9.90	100	0.05	0.05	1.0

Note: The table gives an overview of prices of six old plans (old plan H–M) and the new plan (New plan 2) these six old plans are moved to. Standardized currency.

Figure A.1: Churn rates using actual churn date

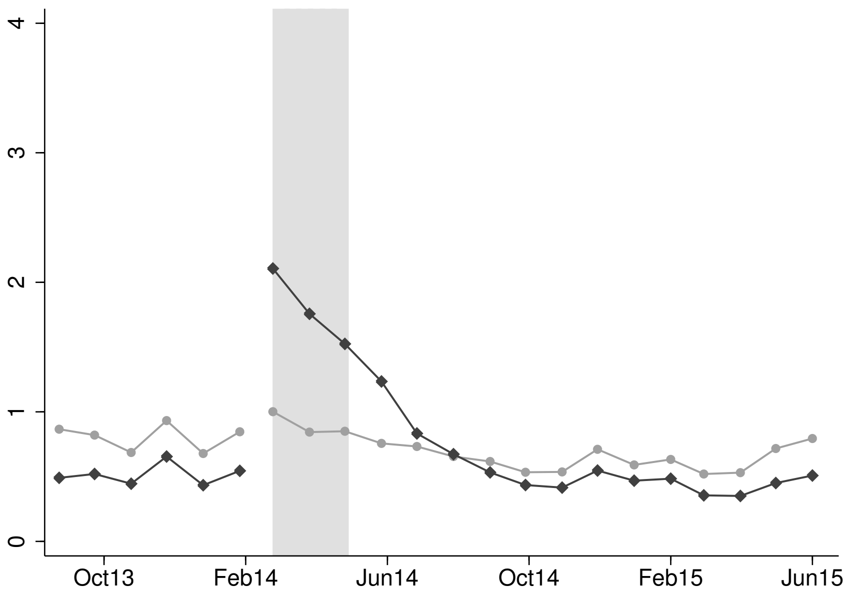

Note: The figure shows the churn rate separately for poked (black line) and non-poked (gray line) consumers. The time from poking to the actual price change is indicated by the shaded area.

Figure A.2: Churn rates within savings categories

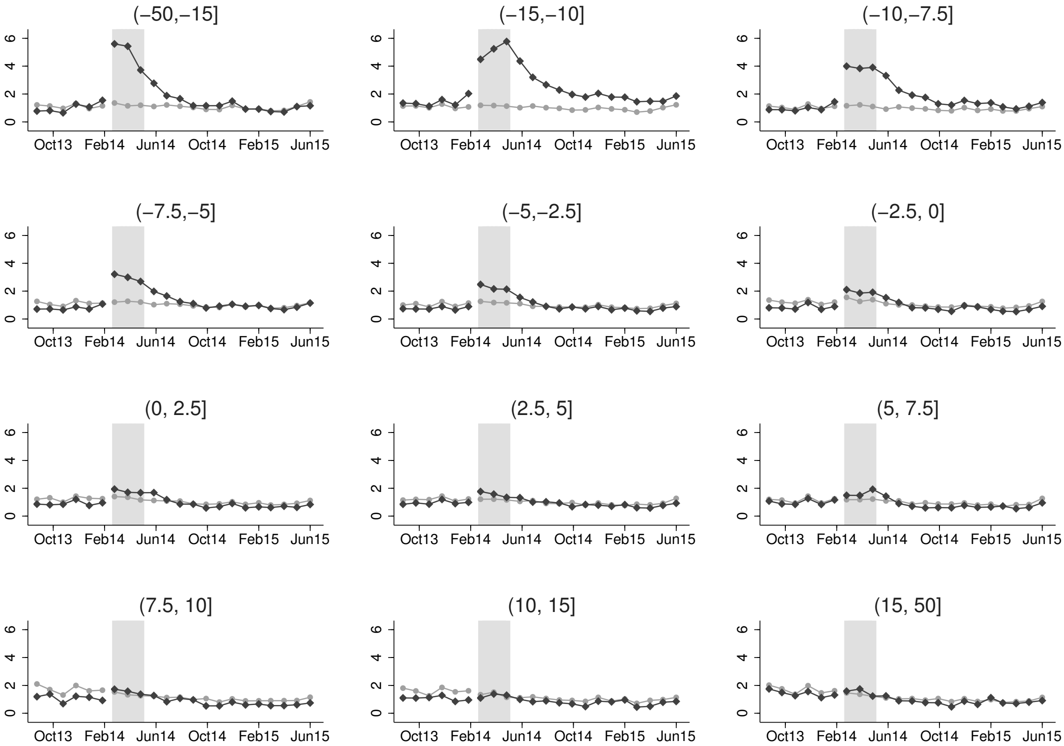

Note: Each panel shows the churn rate for a specific category of savings, separately for consumers who are poked (black line) and all other consumers (gray line). The time from poking to the actual price change is indicated by the shaded area.

Figure A.3: Timing of subscriptions to contracts used to construct restricted sample

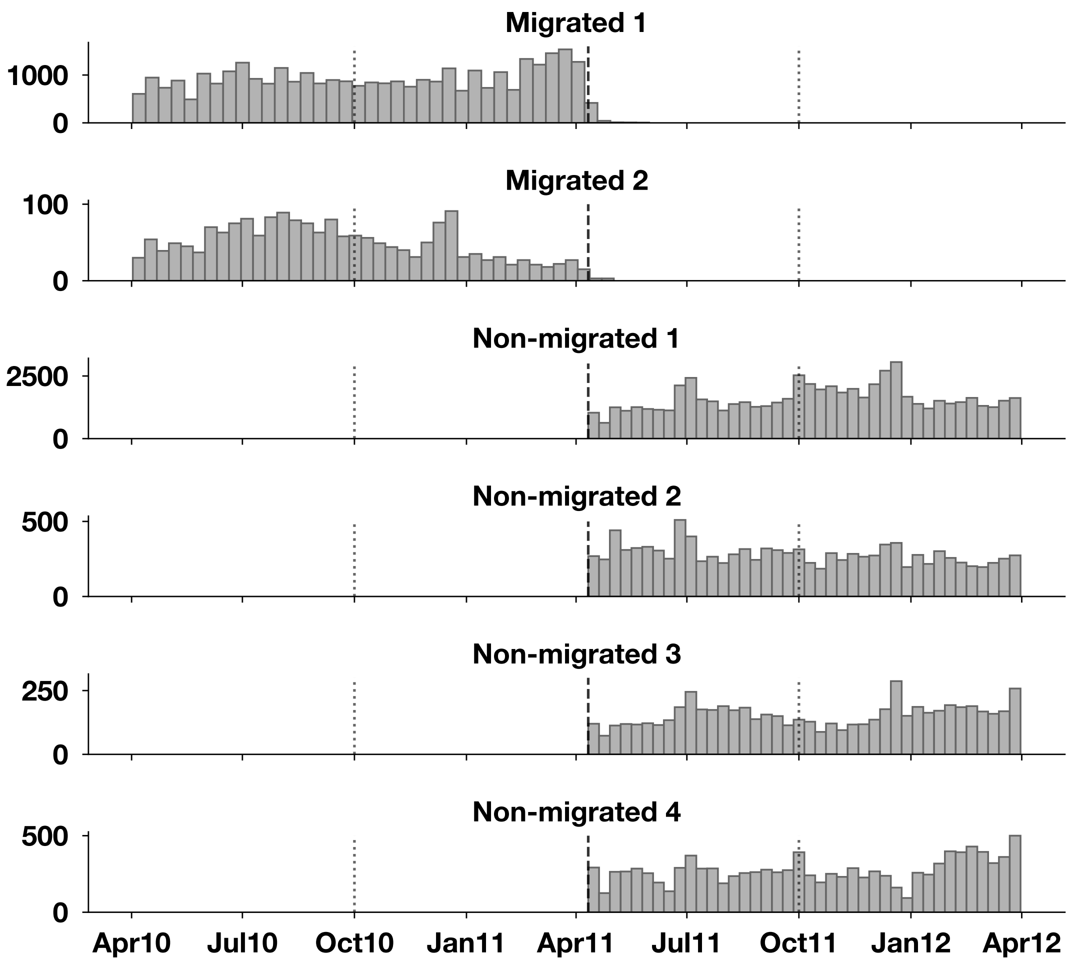

Note: The figure shows the distribution of sales over time around a period where 2 subsequently migrated plans were discontinued and 4 non-migrated plans were introduced. The light, dotted lines indicates the window we use to construct our restricted sample (from October 2010 to October 2011).

Figure A.4: Comparison of churn rates in restricted and full sample

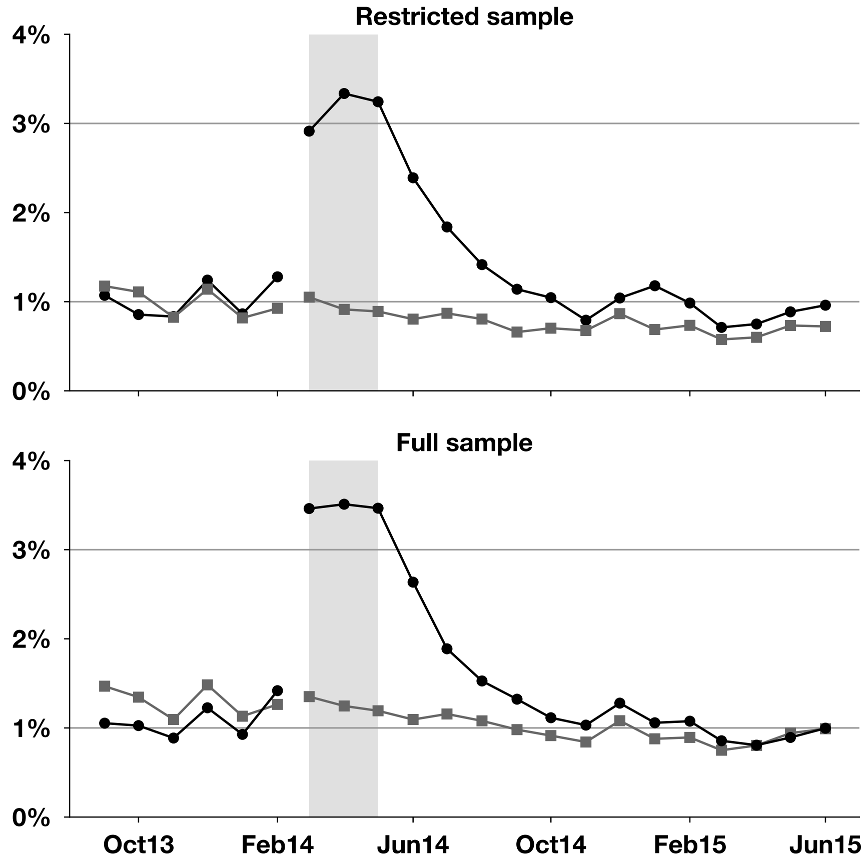

Note: The upper panel (Restricted sample) shows churn rates for the subsample of consumers who signed up for contracts in the 6 months before and after 2 migrated plans were discontinued and 4 non-migrated plans were introduced. The bottom panel (Full sample) shows churn rates for the full sample (equivalent to Figure 2). Churn rates are separated for consumers who are poked (black line) and all other consumers (gray line). The time from poking to the actual price change is indicated by the shaded area.

Figure A.5: Churn rates within age categories

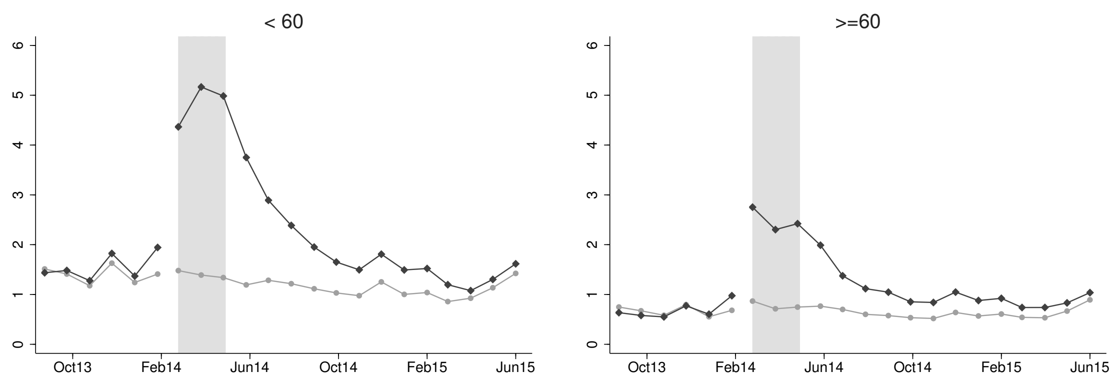

Note: The figures shows churn rates for consumers above 60 years (left panel) and below 60 years (right panel). Churn rates are separated for consumers who are poked (black line) and all other consumers (gray line). The time from poking to the actual price change is indicated by the shaded area.

Table A.3: Overview of main competitor prices at time of poking

Plan	**Prices**					
	Fixed fee	Min. price above incl.	SMS price above incl.	Incl. MB	Incl. min	Incl. SMS
Comp. 1: plan A	20	0	0	1000	unlimited	unlimited
Comp. 1: plan B	30	0	0	5000	unlimited	unlimited
Comp. 1: plan C	10	0.04	0.04	150	150	150
Comp. 1: plan D	40	0	0	8000	unlimited	unlimited
Comp. 2: plan A	30	0	0	3000	unlimited	unlimited
Comp. 2: plan B	20	0	0	500	unlimited	unlimited
Comp. 2: plan C	40	0	0	6000	unlimited	unlimited
Comp. 2: plan D	50	0	0	8000	unlimited	unlimited
Comp. 3: plan A	5	0.04	0.03	0	0	0
Comp. 3: plan B	20	0	0	1000	unlimited	unlimited
Comp. 3: plan C	30	0	0	5000	unlimited	unlimited
Comp. 3: plan D	10	0.04	0.04	150	150	150
Comp. 4: plan A	20	0	0	1000	unlimited	unlimited
Comp. 4: plan B	30	0	0	5000	unlimited	unlimited
Comp. 5: plan A	5	0.04	0.04	0	0	0
Comp. 5: plan B	40	0	0	6000	unlimited	unlimited
Comp. 6: plan A	30	0	0	3000	unlimited	unlimited
Comp. 6: plan B	25	0	0	1000	unlimited	unlimited
Comp. 6: plan C	15	0.05	0.05	200	200	200

Note: The table provides an overview of the main offers by the main competitors at the time of poking (when consumers were informed about the price change). The main competitors do not make any major price changes in the next 3-month period.

Table A.4: Results controlling for save categories

Predicted savings categories	(1)	(2)	(3)
Save (-50,-15]	3.988***	3.990***	3.390***
	(0.178)	(0.177)	(0.149)
Save (-15,-10]	3.771***	3.771***	3.522***
	(0.120)	(0.118)	(0.093)
Save (-10,-7.5]	2.941***	2.916***	2.802***
	(0.133)	(0.131)	(0.117)
Save (-7.5,-5]	2.131***	2.103***	2.096***
	(0.121)	(0.120)	(0.106)
Save (-5,-2.5]	1.385***	1.362***	1.380***
	(0.104)	(0.103)	(0.093)
Save (-2.5,0]	1.008***	0.987***	0.994***
	(0.117)	(0.116)	(0.107)
Save (0,2.5]	0.882***	0.867***	0.881***
	(0.117)	(0.115)	(0.109)
Save (2.5,5]	0.731***	0.712***	0.726***
	(0.130)	(0.129)	(0.123)
Save (5,7.5]	0.632***	0.630***	0.635***
	(0.141)	(0.140)	(0.136)
Save (7.5,10]	0.917***	0.898***	0.932***
	(0.235)	(0.233)	(0.224)
Save (10,15]	0.507***	0.494***	0.541***
	(0.171)	(0.169)	(0.163)
Save (15,50]	0.525**	0.510**	0.521**
	(0.216)	(0.213)	(0.204)
Demographics	No	Yes	Yes
Demogr. interact with treatment	No	No	Yes

*Note: The table shows estimates of the churn increase in the during period per predicted savings category, i.e., γ_3 in Equation 1. Model (1) reports estimates without any control variables. In model (2) we add controls for age, gender, and place of residence, whereas in model (3) we also interact the demographics with treatment (being poked and time period). Robust standard errors clustered on consumers' choice of price plan before the poke, i.e., a combination of price plan and tenure on that plan, 3,089 clusters. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Table A.5: Survey descriptives

		Percent
Female		48.34
Income	Below 20.000	5.10
	20.000-29.999	7.54
	30.000-39.999	14.86
	40.000-49.999	21.95
	50.000-49.999	15.52
	60.000-49.999	8.86
	70.000-49.999	7.54
	80.000-99.999	7.32
	100.000 or above	4.21
	Don't want to answer	7.10
Age	Below 30	9.31
	30-44	17.96
	45-59	27.49
	60+	45.23

Note: The table shows descriptive statistics of gender, income, and age for the survey respondents.

Table A.6: Model coefficient estimates

	Coeff.	Std. err.
Stage 1: Attention parameters		
Constant	−1.87	0.01
Feb14	−0.10	0.01
Mar14	−0.02	0.01
Apr14	−0.06	0.01
May14	−0.05	0.01
Poked	−1.12	0.01
× Feb14	0.01	0.01
× Mar14	0.12	0.01
× Apr14	0.76	0.01
× May14	0.54	0.01
Female	−0.05	0.00
Age		
21–30	−0.10	0.01
31–40	−0.16	0.01
41–50	−0.24	0.01
51–60	−0.34	0.01
61–70	−0.44	0.01
71–80	−0.58	0.01
Residency		
Four largest	0.01	0.01
Urban	−0.01	0.01
Village	−0.06	0.01
Unknown	0.65	0.01
Stage 2: Choice parameters		
Cost	−0.64	0.00
Current provider	0.84	0.01
Other provider	0.22	0.01
Individuals		621,856
Observations		3,019,439
$\ln \mathcal{L}$		−1,086,197

Note: The upper part of the table shows estimated coefficients for Stage 1of the choice model (attention), while the lower part of the table shows estimated coefficients for Stage 2 (choice of plan conditional on being attentive). Current provider refers to the relevant company.

Appendix B Dynamic model and choice probabilities

We can define the ex ante (integrated) value function as

$$\begin{aligned}\bar{V}(j,\mathcal{J}) =& \mathbb{E}\, V(j,\mathcal{J},a,\epsilon) = \sum_a \int V(j,\mathcal{J},a,\epsilon) p_a(a) g_\epsilon(\epsilon) \mathrm{d}\epsilon \\ =& p^A \ln\left(\sum_{j'\in\mathcal{J}} \exp\left\{u(j') + \beta\bar{V}(j',\mathcal{J}')\right\}\right) \\ &+ (1-p^A)\left\{u(j) + \beta\bar{V}(j,\mathcal{J}')\right\}, \end{aligned} \tag{6}$$

where $p^A \equiv p_a(1)$, i.e., the probability of being attentive. The log of summed exponentials is the well-known expression for expected utility when choosing between several options with additive extreme value, type-1 random utility terms, also known as the *inclusive value*. Let $I_\mathcal{J}$ denote the inclusive value from choice set $\mathcal{J}$. When the consumer does not expect any changes in the set of available plans, i.e., $\mathcal{J} = \mathcal{J}'$, we can express a partial solution to the ex ante Bellman equation as

$$\bar{V}(j,\mathcal{J}) = \frac{p^A}{1-\beta(1-p^A)} I_\mathcal{J} + \frac{1-p^A}{1-\beta(1-p^A)} u(j), \tag{7}$$

where the common denominator reflects the possibility of staying inattentive in any number of consecutive periods adjusted for discounting.[39] Note that only the second term depends on the endogenous state variable j, where j reflects the choice made in the last attentive period. When the consumer becomes attentive, the previous choice is not important to the consumer, only the future possibility of repeating the current choice because of inattention.

In periods where the consumer is attentive, the value of a choice j net of ϵ is

$$v(j) = u(j) + \beta\bar{V}(j,\mathcal{J}'), \tag{8}$$

which defines the (conditional) choice probabilities through

$$\begin{aligned}\Pr(j|\mathcal{J}) =& \Pr\left\{v(j) + \epsilon(j) \geq v(k) + \epsilon(k),\ \forall k \in \mathcal{J} \setminus j\right\} \\ =& \frac{e^{v(j)}}{\sum_{k\in\mathcal{J}} e^{v(k)}}. \end{aligned} \tag{9}$$

In the case where the consumer does not expect any changes in the choice set, and we set $\beta = 1$, we can insert for Equation (7) in (8), which simplifies to

$$v(j) = \frac{1}{p^A} u(j) + I_\mathcal{J}, \tag{10}$$

where we see that $I_\mathcal{J}$ is common to all choices, and therefore will not affect the choice probabil-

[39]This expression is only valid if p^A is constant over time.

ities:

$$\Pr(j|\mathcal{J}) = \frac{e^{\frac{1}{p^A}u(j)}}{\sum_{k\in\mathcal{J}} e^{\frac{1}{p^A}u(k)}}.$$

Only the current flow utility of each plan is involved in the decision when consumers expect the supply of plans to be stable. The reciprocal of p^A reflects the expected time until the consumer is attentive again, i.e., the expected time being stuck with the chosen plan without being able to reoptimize.[40]

For consumers who are informed that their current plan will change in the future and no longer be available, the decision problem is different from that outlined above. Though we allow this event to have an effect on choices through its impact on attention, it should be expected to change current behavior even when keeping p^A fixed, since consumers will account for the possibility that they might be inattentive in the future (when the terms actually change). This could lead consumers to change their plan when attentive *prior* to the change in terms, even if they are on the optimal plan considering the current choice set in isolation.

Denote by $\mathcal{J}_0$ the choice set prior to the change, $\mathcal{J}_1$ the choice set after the change, and $\mathcal{J}_C$ the set of plans that will be changed. In the period when the change is implemented, any previous choice $j \in \mathcal{J}_C$ will be changed to $\tilde{j}$. Letting $C(j)$ be an indicator for $j \in \mathcal{J}_C$, and assuming that consumers do not expect any further changes to the choice set in the future after the change is implemented, the ex ante value function in the period in which the change takes place can be written as

$$\bar{V}(j,\mathcal{J}_1) = I_{\mathcal{J}_1} + \frac{1-p^A}{p^A}\left[C(j)u(\tilde{j}) + (1-C(j))\,u(j)\right], \tag{11}$$

where we have set $\beta = 1$. The choice-specific value function in the period before the change ($t_c - 1$) is then obtained by inserting into Equation (8):

$$v_{t_C-1}(j) = I_{\mathcal{J}_1} + C(j)\left(u(j) + \frac{1-p^A}{p^A}u(\tilde{j})\right) + (1-C(j))\,\frac{1}{p^A}u(j),$$

which, for the plans that will be changed next period ($C(j) = 1$), reflects the expected time of being stuck with payoff $u(\tilde{j})$ due to inattention, adjusted for the difference between $u(j)$ and $u(\tilde{j})$ in the same period. To see the intuition behind this expression, imagine a consumer for whom j is the optimal plan in $\mathcal{J}_0$, while there are better plans than $\tilde{j}$ in $\mathcal{J}_1$. Then, there is a trade-off between the gain from getting one last period on the optimal plan against the loss of an expected $\frac{1-p^A}{p^A}$ subsequent periods on a suboptimal plan, where this gain and loss is evaluated against the best plan other than j that will be available both before and after the change.[41] For the plans that will not be changed, the choice-specific value is the same as that in Equation (10). Choice probabilities are given by the multinomial logit formula with $v_{t_C-1}(j)$ for each plan in $\mathcal{J}_0$ as arguments.

[40]Note that $\beta = 1$ implies that $I_{\mathcal{J}}$ is infinite, and should therefore be interpreted as a notationally convenient and reasonable approximation.

[41]We have assumed here that the consumer cannot choose a plan that will be changed in the subsequent period other than the one they are currently subscribed to, which was the case in our empirical setting.

The ex ante value function for the period before the change (which is the relevant continuation value for the decision two periods before the change) can be found by inserting for (11) in (6), which yields

$$\bar{V}_{t_C-1}(j,\mathcal{J}_0) = p^A I_{t_C-1} + (1-p^A) I_{\mathcal{J}_1}$$
$$+(1-p^A)\left[C(j)\left(u(j)+\frac{1-p^A}{p^A}u(\tilde{j})\right)+(1-C(j))\,\frac{1}{p^A}u(j)\right].$$

Again, inserting this value into Equation (8) yields the choice-specific value function two periods before the change:

$$\begin{aligned} v_{t_C-2}(j) =& p^A I_{t_C-1} + (1-p^A) I_{\mathcal{J}_1} \\ &+C(j)\left((2-p^A)u(j)+\frac{(1-p^A)^2}{p^A}u(\tilde{j})\right)+(1-C(j))\,\frac{1}{p^A}u(j). \end{aligned}$$

The general form of the choice-specific value function with a known change s periods in the future is

$$\begin{aligned} v_{t_C-s}(j) =& \sum_{\tau=1}^{s-1}(1-p^A)^{s-1-\tau} p^A I_{t_C-\tau} + (1-p^A)^{s-1} I_{\mathcal{J}_1} \\ &+C(j)\left(\sum_{\tau=0}^{s-1}(1-p^A)^{\tau}u(j)+\frac{(1-p^A)^s}{p^A}u(\tilde{j})\right)+(1-C(j))\,\frac{1}{p^A}u(j), \end{aligned}$$

though we will only need the values up to two periods in advance, since the relevant consumers are notified 2 months prior to the change.

Note that the inclusive value terms will be irrelevant in the choice probabilities, as they shift all choice values by the same amount given the consumer's available choice set.

Issued in the series Discussion Papers 2017

2017

01/17 January, **Agnar Sandmo**, “Should the marginal tax rate be negative? Ragnar Frisch on the socially optimal amount of work”

02/17 February, **Luca Picariello**, “Organizational Design with Portable Skills”

03/17 March, **Kurt R. Brekke**, Tor Helge Holmås, Karin Monstad og Odd Rune Straume, “Competition and physician behaviour: Does the competitive environment affect the propensity to issue sickness certificates?”

04/17 March, **Mathias Ekström**, “Seasonal Social Preferences”.

05/17 April, Orazio Attanasio, Agnes Kovacs, and **Krisztina Molnar**: “Euler Equations, Subjective Expectations and Income Shocks”

06/17 April, **Alexander W. Cappelen**, Karl Ove Moene, Siv-Elisabeth Skjelbred, and **Bertil Tungodden**, “The merit primacy effect”

07/17 May, **Jan Tore Klovland**, “Navigating through torpedo attacks and enemy raiders: Merchant shipping and freight rates during World War I”

08/17 May, **Alexander W. Cappelen**, Gary Charness, **Mathias Ekström**, Uri Gneezy, and **Bertil Tungodden: “**Exercise Improves Academic Performance”

09/17 June, **Astrid Kunze**, “The gender wage gap in developed countries”

10/17 June, **Kristina M. Bott, Alexander W. Cappelen, Erik Ø. Sørensen** and **Bertil Tungodden**, “You’ve got mail: A randomized field experiment on tax evasion”

11/17 August, Marco Pagano and **Luca Picariello**, “Talent Discovery, Layoff Risk and Unemployment Insurance”

12/17 August, Ingrid Kristine Folgerø, **Torfinn Harding** and Benjamin S. Westby, «Going fast or going green? Evidence from environmental speed limits in Norway”

13/17 August, **Chang-Koo Chi**, Pauli Murto, and Juuso Välimäki, “All-pay auctions with affiliated values”

14/17 August, **Helge Sandvig Thorsen**, "The effect of school consolidation on student achievement".

15/17 September, Arild Sæther, "Samuel Pufendorf and Ludvig Holberg on Political Economy".

16/17 September**, Chang-Koo Chi,** Pauli Murto, and Juuso Välimäki**, "**War of attrition with affiliated values".

17/17 September, **Aline Bütikofer** and Giovanni Peri, "The Effects of Cognitive and Noncognitive Skills on Migration Decisions"

18/17 October, **Øivind Schøyen**, "What limits the powerful in imposing the morality of their authority?

19/17 October, **Charlotte Ringdal** and **Ingrid Hoem Sjursen**, "Household bargaining and spending on children: Experimental evidence from Tanzania"

20/17 December, **Fred Schroyen** and Karl Ove Aarbu, "Attitudes towards large income risk in welfare states: an international comparison"

21/17 December, **Alexander W. Cappelen, Ranveig Falch**, and **Bertil Tungodden**, "The Boy Crisis: Experimental Evidence on the Acceptance of Males Falling Behind

2018

01/18 January, Øystein Foros, Mai Nguyen-Ones, and **Frode Steen**, "Evidence on consumer behavior and firm performance in gasoline retailing"

02/18 January, **Agnar Sandmo**, "A fundamental externality in the Labour Market? Ragnar Frisch on the socially optimal amount of work"

03/18 February, Pierre Dubois and **Morten Sæthre, "**On the Effect of Parallel Trade on Manufacturers' and Retailers' Profits in the Pharmaceutical Sector"

04/18 March, **Aline Bütikofer**, Julie Riise, and Meghan Skira, "The Impact of Paid Maternity Leave on Maternal Health"

05/18 March, **Kjetil Bjorvatn** and **Bertil Tungodden**, "Empowering the disabled through savings groups: Experimental evidence from Uganda"

06/18 April, Mai Nguyen-Ones and **Frode Steen,** "Measuring Market Power in Gasoline Retailing: A Market- or Station Phenomenon?

07/18 April, **Chang Koo Chi** and Trond Olsen, "Relational incentive contracts and performance"

08/18 April, **Björn Bartling**, **Alexander W. Cappelen, Mathias Ekström, Erik Ø. Sørensen,** and **Bertil Tungodden**, «Fairness in Winner-Take-All Markets»

09/18 April, **Aline Bütikofer, Sissel Jensen**, and **Kjell G. Salvanes**, «The Role of Parenthood on the Gender Gap among Top Earners»

10/18 May, **Mathias Ekström**, "The (un)compromise effect"

11/18 May, Yilong Xu**, Xiaogeng Xu**, and Steven Tucker**,** «Ambiguity Attitudes in the Loss Domain: Decisions for Self versus Others»

12/18 June, **Øivind A. Nilsen**, Per Marius Pettersen, and Joakim Bratlie, "Time-Dependency in producers' price adjustments: Evidence from micro panel data"

13/18 June, **Øivind A. Nilsen**, Arvid Raknerud, and Diana-Cristina Iancu, "Public R&D support and firms' performance. A Panel Data Study"

14/18 June, Manudeep Bhuller, Gordon B. Dahl, **Katrine V. Løken**, and Magne Mogstad: «Incarceration, Recidivism, and Employment"

15/18 August, Manudeep Bhuller, Gordon B. Dahl, **Katrine V. Løken**, and Magne Mogstad: «Incarceration Spillovers in Criminal and Family Networks"

16/18 August, Pedro Carneiro, Kai Liu, and **Kjell G. Salvanes**: "The Supply of Skill and Endogenous Technical Change: Evidence From a College Expansion Reform"

17/18 August, **Chang Koo Chi**, "An analysis of the two-bidder all-pay auction with common values"

18/18 August**, Alexander W. Cappelen,** Cornelius Cappelen**,** and **Bertil Tungodden, "**Second-best fairness under limited information: The trade-off between false positives and false negatives"

19/18 September, **Aline Bütikofer, Antonio Dalla Zuanna**, and **Kjell G. Salvanes**: "Breaking the Links: Natural Resource Booms and Intergenerational Mobility"

20/18 September, Juan Pablo Atal, José Ignacio Cuesta, and **Morten Sæthre**, "Quality regulation and competition: Evidence from Pharmaceutical Markets"

21/18 October, Orazio Attanasio, Agnes Kovacs, and **Krisztina M**olnar, "Euler Equations, Subjective Expectations and Income Shocks"

22/18 October, Antonio Mele, **Krisztina Molnár**, and Sergio Santoro, "On the perils of stabilizing prices when agents are learning"

23/18 November, Bjørn-Atle Reme, Helene Lie Røhr, and **Morten Sæthre**, "The Poking Effect: Price Changes, Information, and Inertia in the Market for Mobile Subscriptions"

Norges Handelshøyskole

Norwegian School of Economics

NHH
Helleveien 30
NO-5045 Bergen
Norway

Tlf/Tel: +47 55 95 90 00
Faks/Fax: +47 55 95 91 00
nhh.postmottak@nhh.no
www.nhh.no