HOW DO CUSTOMERS LEARN? EVIDENCE
FROM SUBSCRIPTION CONTRACTS

by

I. Ater*
V. Landsman**

Working Paper No 3/2014 November 2014

Research no.: 07840100

* The Leon Recanati Graduate School of Business Administration, Tel Aviv University, Ramat Aviv, Tel Aviv, 69978, Israel. Email: ater@post.tau.ac.il

** The Leon Recanati Graduate School of Business Administration, Tel Aviv University, Ramat Aviv, Tel Aviv, 69978, Israel and at The School of Economics, Erasmus University Rotterdam, the Netherlands. Email: landsman@post.tau.ac.il

This paper was partially financed by the Henry Crown Institute of Business Research in Israel.

How Do Customers Learn? Evidence from Subscription Contracts

Itai Ater Vardit Landsman

Tel Aviv University Tel Aviv University

November 2014

To investigate mechanisms underlying consumer learning, we analyze panel data of 32,650 owners of checking accounts facing a newly introduced menu of three-part tariff subscription contracts. We focus on contract switching decisions as an indication for learning, and identify two triggers for switching: adoption of a contract with an excessively large allowance, and the experience of overage fees. The contract changes occurred quickly after trigger realization, and each led to a different outcome. Customers who experienced overage fees switched to contracts with larger allowances and eventually paid *more* after the switch, while customers who adopted contracts with excessively large allowances switched to contracts with smaller allowances and ended up paying less. We argue that our findings are best explained by directional learning and are inconsistent with Bayesian learning.

Key words: learning, directional, reinforcement, Bayesian, non-linear pricing

JEL Classifications: D01; D03; D81; D83; G21

* Itai Ater is an Assistant Professor of Economics and Strategy at the Faculty of Management at Tel-Aviv University, Israel, Ater@post.tau.ac.il. Vardit Landsman is an Assistant Professor of Marketing at the Faculty of Management at Tel-Aviv University, Israel, and at The School of Economics, Erasmus University Rotterdam, the Netherlands, landsman@post.tau.ac.il.

We are grateful for the comments and suggestions made by Kfir Eliaz, Ido Erev, Rani Spiegler and seminar participants at Tel Aviv University, Hebrew University and Bar Ilan University. We also thank Emanuel Marcu for excellent research assistance. Remaining errors are our own. The project was supported by the Henry Crown Institute of Business Research in Israel.

I. Introduction

The capacity to process new information and learn has long been acknowledged as a fundamental component of decision making and behavior, in economics and in general (Simon 1959, Tversky and Kahenman 1971, 1973). For instance, theorists assume that information processing enables individuals to reach equilibrium behavior (Smith 1982; Sobel 2000). Labor economists study how firms use new information to learn about productivity and wage dynamics (Farber and Gibbons 1996), and macroeconomists study how individuals form expectations about future inflation (Woodford 2003).

Bayesian updating, combined with expected utility theory, is by far the most common approach used in economics to model learning. However, early experimental studies (e.g., Kahneman and Tversky 1972; Tversky and Kahneman 1973) have demonstrated that individual behavior often violates the underlying assumptions of this approach. Experimental studies have further shown that reinforcement learning, a learning theory developed in psychology, can better predict individuals' behavior (Roth and Erev 1995; Erev and Roth 1998). A third theory, directional learning, developed by Reinhard Selten, has also been shown to outperform alternative learning theories in many cases (Selten and Stoecker 1986; Selten and Buchta 1999; Selten et al. 2005; Selten and Chmura 2008). Although these studies offer many important insights regarding individual learning, their findings are based exclusively on lab experiments. Thus, important questions remain open. What characterizes individuals' learning behavior in a real market environment? Which learning theory best describes behavior?

Empirically distinguishing among learning theories is challenging. First, the analyzed data should enable individuals' decisions to be tracked over time, while accounting for the nature of the new information that these individuals were exposed to, and the consequences of their decisions. Such data are essential for the identification of the learning mechanisms that are at play, yet are rarely available to researchers. Second, even if such data are available, researchers often do not observe the set of alternatives that an individual faces, which makes it impossible to evaluate customers' alternative choices. Third, alternative learning theories often provide similar predictions, and it is therefore difficult to support one theory and refute another. For instance, most, if not all, learning theories predict that individuals are less likely to repeat unsuccessful choices. Given these challenges, it is not surprising that only a few studies have examined individual learning, and that few empirical studies have investigated which learning theory best explains individual learning behavior. In particular, we are not aware of previous field studies that explore the extent to which directional learning theory actually explains the dynamic behavior of individuals.

We address this gap in the literature by taking advantage of a rich set of panel data on decisions made by individual customers with regard to a menu of subscription contracts. Subscription contracts provide a natural setting in which to study individual learning, because contract choice and actual usage are separated in time, and because customers typically do not forecast their future needs perfectly at the time of subscription choice. Specifically, we study decisions made by a group of 32,650 holders of checking accounts in a large retail bank in an OECD member country, who each chose a subscription contract from a newly introduced menu of three-part tariff subscription contracts. Three-part tariff contracts consist of a *fixed fee*, an included *allowance* of units for which the marginal price is zero, and an *overage* payment—a positive marginal price for additional usage beyond the designated allowance. Such pricing schemes are commonly used in the cellular, Internet and car leasing industries. Bank customers could select a contract from the new menu or by default continue to be charged according to the pay-per-use pricing scheme that was in place prior to the menu's introduction. We observe the usage and monthly commission payments of the customers we investigate, before and after choosing one of the new contracts. In particular, we track their behavior over 30 months (including 6 months before the new menu was introduced), and observe their initial contract choices and possible contract switches. The data also include rich information on each account holder, including age, number of children, salary, account tenure, loans and savings.

We use our unique setting to test alternative learning theories. To motivate our empirical analysis, in Section 3 we describe main features of Bayesian, reinforcement and directional learning theories. We use this description to derive testable hypotheses that we later examine in the data. In the analysis, we emphasize instances in which the predictions of the alternative learning theories clash.

Our starting point for the empirical analysis is a customer's initial choice of a three-part tariff contract, referred to as the contract adoption decision. We show that 69 percent of the customers in our data set adopted contracts with allowances that were larger than the allowances of the contracts that would have minimized their costs, given their usage of the service (either before or after contract adoption).

Our main analysis focuses on customers' decisions to update their initial contract choices, which we consider as an indication of learning. In our analysis, we use panel data analysis and duration models to study the *determinants*, the *speed*, and the *consequences* of a customer's decision to switch to an alternative three-part tariff contract or to quit, i.e., return to the 'old' pay-per-use pricing scheme after adopting one of the three-part tariff contracts. Notably, the contract options are presented to the prospective customer in order of the size of the allowance or fixed payments, making our setting attractive for testing directional learning theory. Overall, we observe 2,268 contract switches in our data, and 2,030 customers who returned to the old pay-per-use pricing scheme.

We document the following findings. First, we identify two factors that trigger contract switching: overage payments, incurred when a customer exceeds his or her contract allowance, and adoption quality, measured by the ratio between post- and pre-adoption payments. Second, the switching decision associated with a given trigger takes place shortly after the trigger is realized. Third, we observe that the two identified triggers have opposing effects on the direction of the subsequent contract switch. Overage payments are associated with a switch to a 'higher' contract (i.e., a contract with higher fixed fees and larger allowances), while low adoption quality is typically associated with a switch to a 'lower' contract (i.e., a contract with lower fixed fees and smaller allowances). Fourth, we find that customers who choose to switch to lower contracts decrease their overall commission payments (a commission payment is defined as a customer's total monthly payment to the bank; among contract adopters, the commission is made up of the fixed fee and the overage fees). On the other hand, and importantly, customers who switch to higher contracts end up paying more after the switch. Finally, we also analyze decisions to return to the old pricing scheme (quit) and find that quitting is positively associated with both overage payment and poor adoption quality (i.e., adoption of a contract that increases the customer's monthly commission payments). We rely on the predictions developed in Section 3 to argue that our findings are inconsistent with the standard Bayesian learning model, and that directional learning theory has better predictive power than reinforcement learning. In particular, it is difficult to explain, on the basis of standard Bayesian models, why individuals who experience overage payments tend to switch to contracts with larger allowances and pay *more* after the switch.

The rest of the paper is organized as follows. In Section 2 we review the related literature. In Section 3 we describe the main features of the learning theories we consider and develop testable hypotheses. In Section 4 we describe the industry we study and the data we analyze. Section 5 contains the empirical analysis of adoption, contract switching and quitting decisions. In Section 6, we discuss and offer concluding remarks.

II. Related literature

Only a few studies have used micro-level panel data to examine learning patterns (e.g., Ito 2014; Grubb and Osborne 2014); this scarcity is most likely due to data limitations. These studies, which generally found that learning improves individuals' economic outcomes, can be broadly divided into two categories. The first category, like this paper, investigates individuals' decisions with regard to a *menu* of contracts and typically compares customers' payoffs and usage patterns before and after contract switching (Miravete 2003; Agarwal et al. 2006; DellaVigna and Malmendier 2006; Ketcham et al. 2012). Ketcham et al. (2012) find that 81 percent of a sample of Medicare part D plan enrollees reduced their payments between 2006 and 2007 by switching plans, with an average reduction in overspending of 298 dollars. Miravete (2003) tracks subscribers of calling plans 3 months before and 3 months after their subscription decisions and finds that switchers made lower payments after switching. Unlike our study, these studies show that subscribers reduce their payments after they switch plans.

The second category of studies investigates individual usage of a service under a given contract before and after an 'unexpected' event, such as high payments. Haselhuhn et al. (2012) analyze data from Blockbuster and show that a customer who experiences paying a late-return fee is more likely to return his or her next rented DVD on time. Agarwal et al. (2013) show that credit card holders who incur add-on fees in a given month are unlikely to incur these payments again in the following months. Like our work, these studies consider unexpected events as a trigger for learning.

Few studies in the finance literature have used individual panel data to examine the impact of personal experiences on financial decisions (Kaustia and Knupfer 2008; Chiang et al. 2011; Choi et al. 2009). These papers did not examine the applicability of directional learning and have generally found evidence that individual behavior is consistent with reinforcement learning. Finally, Malmendier and Nagel (2011) use repeated cross-sectional data on household asset allocation and find that personal experience with macroeconomic shocks affects their financial risk-taking. They also find evidence for the recency effect, that a trigger—a macroeconomic shock in their analysis—has a short-term impact on households' financial choices.

Our findings are also related to theoretical studies on non-linear pricing. Recent theoretical studies have shown that firms may adopt non-linear pricing schemes to take advantage of consumers' sensitivity to the marginal price they expect to pay (Herweg and Mierendorff 2013; Bordalo et al. 2013). Indeed, our findings also suggest that consumers' contract choices are sensitive to the marginal price – overage payments – they pay. In contrast to our findings, Ito (2014), who also empirically investigates a setting of non-linear pricing, finds that consumers' electricity usage is influenced by average prices rather than by marginal prices. One possible explanation for this difference is customers' ability to respond to the marginal price by switching to alternate contracts. Unlike our setting, in Ito (2014), customers could only respond to the marginal price by adjusting their usage level yet could not adjust by switching to an alternative contract.

Finally, our paper also adds to the behavioral economics literature. Although there is growing evidence that customers exhibit different biases that cause them to diverge from rational behavior (e.g. DellaVigna, 2009), it is possible that experience may reduce or eliminate such biases, which would suggest that their practical importance is ultimately limited. Indeed, previous field studies on experience have shown that individuals tend to overcome initial biases after they gain experience (e.g. List 2003, 2004). However, Ater and Landsman (2013) show that choice biases do not disappear as customers gain experience. These mixed findings reinforce Levitt and List's (2008) call for further investigation of "how markets and market experience influence behavior." A better understanding of the learning mechanisms underlying consumer behavior could elucidate which biases are likely to be eliminated with experience and which are likely to remain.

III. Learning Theories and Testable Hypotheses

In this section, we first review basic features of Bayesian updating, reinforcement learning, and directional learning. We then emphasize the main differences among the three theories and, on the basis of these differences, develop testable predictions that we later examine in the data.

III.1. Learning Theories

III.1.1 Bayesian updating

Also referred to as belief updating, or Bayesian learning, Bayesian updating is the prevalent approach used in economics to model individual learning. This approach assumes that individuals have some prior belief regarding the desirability of different alternatives. They then gain information over time regarding this desirability through various signals, and this information enables them to update their initial beliefs, using Bayes' rule. As a normative approach, Bayesian updating—in combination with expected utility theory—implies that individuals improve or, at least, do not impair their well-being as they gain more information. In our context, this implies that once a customer learns about the fit between alternative contracts and his actual usage, he will not switch to a contract that is expected to increase his commission payments.[1]

Kahneman and Tversky (1972; Tversky and Kahneman 1973) were among the first to demonstrate that individuals systematically violate the standard assumptions of Bayesian inference. Since then, several other studies have found experimental evidence for violations of the Bayes rule (e.g., Ortoleva 2012 and the references therein). Following these studies, researchers have suggested theoretical modifications, often within the standard Bayes inference framework, that can accommodate individuals' underlying behavior (e.g., Rabin 2002; Gennaioli and Schliefer 2010). Rabin (2002), for instance, develops a model that accommodates the 'law of small numbers' bias. In this model, individuals over-evaluate the likelihood that a small sample resembles the parent population from which it is drawn.

Outside the Bayesian inference framework, researchers have developed alternative learning theories that rely on insights from psychology and do not necessarily follow the standard assumptions of economic optimization (see Harstad and Selten 2013; and Crawford 2013 for insightful discussions). A classic example of such boundedly-rational learning models is the theory of reinforcement learning.

III.1.2 Reinforcement learning

The core component of reinforcement learning, a theory originally developed by psychologists, is the law of effect (Thorndike 1932). This law implies that personal experiences strengthen or weaken the propensity of individuals who have taken a given action to engage in

[1] Given the relatively small stakes involved in our setting, risk aversion cannot explain why individuals choose contracts that do no minimize their payments (Rabin 2000).

that action again (Erev and Roth 1998). In other words, the payoff yielded by a given choice in a preceding period determines the increase or decrease in its choice probability in the following period. Higher payoffs are associated with higher future choice probabilities, whereas lower payoffs are associated with lower future choice probabilities. Reinforcement theory does not take into account individuals' beliefs about what other options would have yielded (i.e., forgone payoffs). Rather, it assumes that the decision maker considers only the payoffs yielded by his or her own past choices. A basic implication of reinforcement learning is that after going through a negative experience, individuals will tend to switch to another alternative. Notably, however, the theory of reinforcement learning does not indicate which alternative is more likely to be adopted.

III.1.3 Directional learning theory

Directional learning (or learning direction) theory (Selten and Stoecker 1986; Selten and Buchta 1999; Selten et al. 2005) is a qualitative theory about learning in repetitive decision tasks. The following simple example is often used to introduce the basic principle of this theory. Consider an archer who tries to hit the trunk of a tree. If the arrow misses the tree on the left side, the archer will tend to aim more to the right in the next round, and in the case of a miss to the right the following aim will be more to the left. The behavior of the archer is based on a qualitative causal picture of the world, where the directional change of the aim 'offsets' the deviation of the arrow from the trunk in the previous round. Thus, a decision maker evaluates what she could have done better last time and adjusts the decision in this direction.

In contrast to reinforcement learning, in which experienced payoffs are the only factor influencing subsequent decision making, in directional learning the additional payoff that might have been gained through other actions is a key component of the learning process. According to directional learning theory, the comparison of experienced payoffs with hypothetical payoffs guides the decision maker. Since counterfactual causal reasoning about the past is a crucial feature of directional learning theory, a negative experience not only is likely to lead to abandonment of the choice that resulted in that experience, but is also likely to lead to a correction in a direction that reduces the likelihood of re-experiencing the negative experience.

Despite the fundamental difference between them, reinforcement and directional learning have several important features in common. In particular, both theories emphasize the recency effect (Hogarth and Einhorn, 1992; Hertwig et al. 2004; Ockenfels and Selten 2014). In addition, in both theories individuals assess the 'negativity' of their current experience against a reference point (Erev and Roth 1998, Selten et al. 2005).

III.2. Development of Testable Hypotheses

In the empirical section of this paper we study choices by individuals who face a menu of checking account contracts. Our setting is different from lab experiments, which typically rely on a repetitive structure of multi-period games in which individuals' decisions and their consequences are evaluated in each period. Accordingly, we need to adjust the theoretical predictions to our setting. First, since in practice customers do not necessarily make a decision in every period, we focus only on contract switching decisions. We posit that the customers who switch contracts have gone through an active post-adoption decision-making process, which we consider an indication of learning. Second, we need to identify the triggers for contract changes, and their potential different effects on customers' eventual actions and payoffs.

In our setting the contract menu is ordered according to the size of the allowance, which corresponds to the size of the fixed payments. Thus, our setting is attractive for testing the predictions of directional learning. Accordingly, we define a switch to a contract with a larger allowance/higher fixed payment compared with the customer's pre-switch contract as an *upward* switch, and a switch to a contract with a smaller allowance/lower fixed payment compared with the pre-switch contract as a *downward* switch. On the basis of the characteristics of our setting and the review of the learning theories above, we formulate the following testable hypotheses:

Hypothesis 1. According to Bayesian learning theory, customers who switch contracts will not end up paying more after the switch than they paid before the switch.

Hypothesis 1 is essential for distinguishing between standard Bayesian learning and alternative learning theories, which do not necessarily involve economic improvement. According to the standard Bayesian inference framework, a customer who learns from new informational signals is predicted to act in a manner that does not worsen his economic situation. In our setting this leads to a natural prediction that customers will only switch to contracts that entail lower payments.

Another indication of Bayesian updating is that individuals rely on prior, or base-rate, information when changing their existing contracts (e.g. Tversky and Kahneman 1973; Grether 1980). According to Bodoh-Creed and Rabin (2014), if individuals neglect base-rate information, then we can expect to observe a recency effect. Thus, we predict:

Hypothesis 2. According to Bayesian learning theory, individuals are not subject to the recency effect (i.e., individuals do not underweight prior beliefs and are not more likely to react to recent experiences as compared to more distant experiences).

We now turn to reinforcement and directional learning theories. Both theories require us to consider negative experiences that customers undergo. We focus on two types of negative experiences that serve as potential triggers for subsequent changes. Each of these experiences is associated with a different 'reference payment', which customers use to evaluate their actual monthly commissions. The first negative experience is overage payment, which occurs when a customer exceeds the allowance defined by his or her chosen contract. Overage payments are paid per transaction on top of the fixed contract fee, which, as a dominant feature of the chosen contract, serves as a natural reference payment (Herweg and Mierendorff 2013). Since, by definition, overage payments are additional costs that the customer incurs, we consider them a negative experience.

The second negative experience we consider is poor adoption quality. Adoption quality is calculated as the ratio between the average monthly payments post-adoption and those payments prior to contract adoption. Here, we conjecture that customers use their pre-adoption payment as the reference payment for the assessment of their post-adoption experience. Adoption quality with a value greater than one implies that a customer's mean monthly payment after adoption is higher than his or her average monthly payment before adoption, and the customer is likely to perceive this experience negatively. However, adoption quality can also be positive: a value lower than one implies that contract adoption has led to lower payments, and the customer is likely to perceive this experience positively. Based on this discussion, we put forward the following hypothesis.

Hypothesis 3. According to reinforcement learning theory, customers who undergo negative experiences are expected to switch to alternative contracts or quit their subscription contracts altogether, yet there is no a priori prediction for the specific contracts to which they switch.

According to reinforcement learning, negative experiences lead customers to move away from their current choices. Yet this theory does not enable us to predict which alternative contract will be chosen. Within the directional learning theory framework, however, it is possible to make a more concrete prediction regarding the choices made after negative experiences. In particular, according to directional learning, the contract change will be in a direction that reduces the likelihood of recurrence of the negative experience. More specifically,

Hypothesis 4a. *According to directional learning theory, customers who undergo the negative experience of overage payments will switch upward or quit the new subscription contracts altogether.*

Hypothesis 4b. *According to directional learning theory, customers who experience low adoption quality will switch downward or quit the new subscription contracts altogether.*[2]

Though the standard Bayesian framework cannot explain switches to contracts that lead to increased payments, or behaviors that are consistent with the recency effect, modifications to the standard Bayes framework may be able to accommodate such phenomena. Specifically, according to Rabin (2002), customers may interpret overage payments as a signal for increased usage and may thus overestimate their future usage. If indeed such overestimation of future usage occurs, customers are predicted to switch to contracts with larger allowances. Notably, these customers are not expected to quit their contracts and revert to the old pay-per-use pricing scheme. Clearly, under a pay-per-use scheme, higher usage is associated with higher payments; therefore, under Bayesian learning theory, customers who (erroneously) expect increased future usage are unlikely to choose this type of contract. We thus hypothesize:

[2] As we later show, low quality is closely associated with the adoption of a contract with an excessively large allowance.

Hypothesis 5. According to adjusted Bayesian learning theories, customers may misinterpret the information provided by overage payments as an indication of increased future activity and therefore switch upward, to contracts with larger allowances. However, such customers are not expected to quit their contracts altogether.

IV. Industry Background and Data

IV.1. Economic Environment.

We use data on the introduction of a menu of three-part tariff contracts by a large commercial bank that operates in a developed OECD member country. The bank is one of three large banks that, in the analyzed time period, collectively controlled about 85 percent of the market. Over the years of data collection, relatively few bank customers switched between banks. The introduction of the new pricing scheme that we study followed a public outcry over the complexity of banks' commission structure. The bank from which we obtained the data is a leading bank in the country and was the first to offer the new pricing scheme to its customers. Throughout the paper, we convert the local currency into nominal dollars.

IV.2. Data

IV.2.1. Three-part tariff contracts

The new three-part tariff contracts provided an alternative to an 'old' pricing scheme, which was the system used by all banks operating in the country at the time of the new plans' introduction. Under the old pricing scheme, customers paid a commission (ranging from a few cents to as much as $7) for each transaction they engaged in. Customers who did not choose a new contract following the menu's introduction continued, by default, to use the old pricing scheme (continuing to use the old pricing scheme required no active choice on the customer's part). A customer who adopted a new service contract was free at any time to switch to a different contract or to return to the old scheme ('quit'). To adopt, switch or quit, the customer simply had to call his or her bank branch or the bank's call center; there was no requirement to arrive in person, sign documents, or pay any switching fees.

Each three-part tariff contract entailed a fixed monthly fee, which covered monthly allowances for three types of transactions: check deposits, transactions through direct channels (e.g., Internet or using a touch-tone telephone), and transactions that involve interaction with a clerk at a bank's branch or through a call center.[3] Transactions exceeding these allowances entailed overage payments, paid in addition to the basic contract cost (overage fees of $0.3 for each check deposit or direct channel transaction, and $1.2 for each transaction involving human interaction; fees per transaction were consistent across different contracts). Table 1 presents the details of two three-part tariff contracts: the least expensive contract (contract 1)—the contract with the lowest allowance; and the second most expensive contract (contract 5)—the contract with the next-to-highest allowance.[4] Throughout the analysis, the number of the contract is an indication of the size of the allowance (e.g., contract 2 has a larger allowance than contract 1 and entails a higher fixed payment.) We also use this ordering of contracts to analyze the direction of the contract switching decision ('upward' to a contract with a larger allowance or 'downward' to a contract with a smaller allowance).

IV.2.2. Sample and data

Our data consist of information on 32,650 checking accounts whose holders subscribed to one of the three-part tariff contracts over the sample time period. This list of checking accounts was extracted from an initial list of about one million accounts that the bank had

[3] Three-part tariff contracts for cellular service also typically include three types of allowances: voice, text and data.
[4] In the month when the new contracts were introduced, bank customers could choose from a menu of four three-part tariff contracts. Nine months after the first four contracts were introduced two new contracts were added to the existing set of contracts. After contracts had been offered to customers, they remained available throughout the investigated timeframe, with two exceptions: one contract from the set of the four initial contracts was removed from the choice set nine months after its introduction, and another contracts was altered such that its allowance for direct channels was reduced (customers who chose these contracts before these changes could still use them afterwards).

identified as potential candidates for the service. The initial list of potential accounts was reduced to include only accounts that were active for at least six months at the time that the new service was introduced and that were considered the primary accounts of the account holders. In addition, accounts held by very young customers and accounts for which certain indicators, such as the age or the address of the customer, were missing, were also excluded. To construct the actual sample of accounts we used a layer sampling procedure based on the time of contract adoption. That is, all accounts were ordered according to the date on which the account holder adopted a three-part tariff contract. We then selected every tenth account for the final sample (see Landsman and Givon (2010) for further discussion of the data). The data were collected over the course of 30 months (from 6 months before service introduction until 24 months after introduction).

For each account and for each month, the data set contained the following information: (i) the contract (or contracts) used for the account during that month; (ii) the number of transactions of each type (check deposits, direct-channel transactions and clerk-assisted transactions) carried out by the account holder in that month; (iii) the number of information inquiries performed by the account holder in that month; (iv) additional characteristics, including general characteristics (e.g., account tenure and social security payments deposited into the account), financial characteristics (e.g., income and the monthly levels of savings and loans), and demographic characteristics (e.g., customer age and socio-demographic index[5]). Our data also include the number of direct marketing calls made to each customer to introduce the possibility of choosing from the menu of new contracts. To protect customers' privacy, each account number was encrypted in a way that still enabled us to track that account through

[5] A scale of 1 to 10. Higher values indicate a higher socio-demographic status for the address of the customer.

the entire research data set.[6] Of the 32,650 customers who adopted one of the contracts, 2,268 eventually switched to one of the other three-part tariff contracts, while 2,030 opted to return to the old payment system.

V. Analysis

We start our empirical analysis with contract adoption decisions. We show that customers tend to choose contracts with allowances larger than the allowances of their cost-minimizing contracts. Next, we turn to the main analysis and examine how customers' experience with the new contracts affects subsequent contract switching decisions, and the consequences of these switching decisions. In the final part of this section we also examine the determinants, timing and consequences of quitting decisions.

V.1. Contract Adoption

In our data, 32,650 customers adopted one of the available three-part tariff contracts. We take advantage of the information on the set of contracts available to customers, and on their usage before and after adoption, to examine how customers' actual contract choices compare with their 'optimal' choices, i.e., the choices that would have minimized their costs. Specifically, for each customer we first compute the payments that he or she would have paid under each of the available contracts in a given period. Next, we identify the contract that would have yielded the lowest payment for that customer. Tables 2A and 2B present the distributions of actual versus optimal choices, for each contract, based on account usage 3 months before adoption (i.e., ex-ante approach, Table 2A) or 3 months after adoption (ex-post approach, Table 2B).[7] For example, in Table 2A the number in the second column of the row

[6] Due to confidentiality concerns we are not allowed to reveal summary statistics for the following variables: salary, loans, savings, monthly mean positive balance, and monthly mean negative balance. We use these variables in the regression analysis.
[7] Changing the relevant time-period before or after adoption has little effect on the fraction of customers who chose their cost-minimizing contract. In Appendix A we provide more details on the calculations used for the optimality assessment.

that corresponds to contract 5 indicates that—according to usage during the 3 months before adoption—the optimal contract for 4.6 percent of the adopters of contract 5 is contract 2. The diagonals in Tables 2A and 2B represent the percentage of customers who actually chose their cost-minimizing (optimal) contracts. The vast majority of non-optimal contract choices were for contracts with larger allowances than the allowance offered by the cost-minimizing contract (i.e., there is a large concentration of choices below the diagonals of Tables 2A and 2B). If we aggregate over customers and the different contracts, we find that only 29 percent of contract adopters actually chose the contracts that minimized their payments to the bank.[8] Furthermore, more than 69 percent of the customers adopted contracts with excessively large allowances and higher fixed payments compared with their cost-minimizing contracts. In fact, we find that customers could have reduced their monthly payments, on average, by nearly 30 percent had they chosen their cost-minimizing contracts. This pattern corresponds to a phenomenon known in the literature as flat-rate bias (e.g. Lambrecht et. al 2007).

V.2. Contract Switching

In this section we focus on customers who adopted three-part tariff contracts and later switched to different three-part tariff contracts. Our underlying assumption here is that these switches are indicative of the learning process that customers go through. Throughout the analysis, we distinguish between customers who switched to contracts with lower allowances ('downward-switchers'), and customers who switched to contracts with larger allowances ('upward-switchers'). This classification is central to directional learning theory and is

[8] If customers' account usage is sensitive to the marginal price they pay then our ex-post optimality measure is a lower bound for the real optimality level. Furthermore, when we performed the same calculation, considering also the possibility to remain with the old pricing scheme in the optimality analysis (i.e., we consider the possibility of non-adoption), we find that 17 percent and 18 percent of contract adopters chose the contracts that minimized their payments to the bank (based on their ex-ante and ex-post usage, respectively).

important given customers' strong tendency to initially adopt contracts with larger than optimal allowances.

V.2.1. Determinants and speed of switching

V.2.1.1. Descriptive statistics

Table 3 shows summary statistics for all adopting customers, as well as for the sub-groups of adopting customers who eventually switched contracts, and for adopting customers who eventually returned to the old payment scheme. Customers who adopted three-part tariff contracts paid, on average, $0.43 in monthly overage payments. However, the monthly overage payment varies substantially across customers. As can be seen in columns 2 and 3 in Table 3, customers who eventually switched to contracts with lower fixed fees (downward-switchers) paid on average only $0.21 in overage payments prior to switching, while customers who eventually switched to contracts with higher fixed fees (upward-switchers) paid on average $1.81 in overage payments prior to switching (among switchers, this calculation includes only overage payments prior to switching). In fact, while 91 percent of upward-switchers paid overage payments before switching, only 21 percent of downward-switchers experienced such payments.

Our measure for adoption quality is the ratio of payments after contract adoption (up until switching/quitting if they take place) to payments before contract adoption. Thus, higher values indicate worse adoption quality. Figure 1A plots the distribution of this variable for all adopting customers (mean of 1.17, std. 0.45). Figure 1B plots the distribution of mean adoption quality for switching customers, grouped according to the direction of their switches. We see that customers who switched downward were more likely to exhibit low adoption quality (higher values for the adoption quality measure), compared with customers who switched upward (mean of 1.30 vs. 1.12, for upward- and downward-switchers, respectively). Overall, the descriptive findings seem consistent with the idea that customers respond to the respective triggers by switching to contracts in which the negative events they experienced are less likely to recur (Hypothesis 4a and 4b).

To provide preliminary descriptive evidence that imply that customers make decisions on the basis of recent information, we compare the elapsed time until the switch for downward- and upward-switchers. We expect that contract switches driven by poor adoption choices will occur before contract switches triggered by overage payments since customers are likely to realize that they made an ill-suited adoption choice shortly after adoption. Overage payments, on the other hand, are experienced only after a customer exceeds his or her allowance, not necessarily right after adoption. Figure 2 plots the cumulative distribution of elapsed time to switch, divided into downward- and upward-switchers. As can be seen in the figure, downward switching decisions predominantly take place during the first few periods after contract adoption. Nearly 50 percent of downward switching decisions occur within 4 months after contract adoption. In contrast, upward switching decisions are not concentrated in the months after contract adoption, and only 23 percent of upward switching decisions occur within the first 4 months after the initial contract adoption. Figure 3 further illustrates the strong time proximity between overage payment and upward switching decisions. The figure plots the elapsed number of months between the last overage payment that a customer experienced and the subsequent switching decision. Contract switches occur shortly after overage payments. To further support the descriptive findings we now turn to regression analyses.

V.2.1.2. Identifying the triggers for switching

We now aim to quantify the connection between adoption quality and the direction of subsequent switching decisions (i.e., decisions to switch upward or downward). To this end we estimate the post-adoption change in payments (yet only for pre-switching periods) among all customers, and then separately for upward- and downward-switchers. The idea is to show that, even after controlling for the level of activity (i.e., number of transactions) and other account characteristics, downward-switchers exhibit a larger increase in their post-adoption monthly payments compared with customers who did not eventually switch or customers who subsequently switched upward. We estimate the following panel data fixed-effect regression:

$$(1)\quad Log(payment)_{it} = \beta_0 + \beta_D D_{Adoption\,it} + \beta_A Log(Activity_{it} + 1) + \beta_X Log(X_{it} + 1) + \alpha_i + \eta_t + \varepsilon_{it}$$

where the dummy variable $D_{Adoption,it}$ equals one if customer i has adopted a three-part tariff contract by time t, and zero otherwise. $Activity_{it}$ is a matrix that includes for each account i the number of clerk-assisted and direct transactions, and the number of checks deposited in month t. X_{it} includes account-level characteristics that can vary over time, such as salary amount deposited, number of account owners, number of salaries, social security payments, loans and savings. We implement log-transformation for all the variables that are not binary variables. Finally, we also include account (α_i) and time (η_t) fixed effects to control for unobserved differences across customers and unobserved time trends. The standard errors at the individual account level are clustered.

Column 1 in Table 4 presents the estimation results of Equation (1) for all adopting customers. As shown in the table, following adoption, customers' monthly commissions increased on average by 9.1 percent. In column 2 we focus on all switching customers, and find that, for these customers, commissions increased by only 5.3 percent after the initial contract adoption and before switching. Columns 3 and 4 present the estimation results for customers who subsequently switched upward or downward, respectively. We find that upward-switchers, on average, did not pay more after their initial adoption (before switching) than they had prior to adoption. Downward-switchers, in contrast, did pay considerably more after initial adoption than they had paid prior to adoption. Specifically, these customers experienced an increase of 16.9% in monthly commissions compared with pre-adoption months, after controlling for their level of activity. This latter finding suggests that low-quality adoption decisions triggered downward switching decisions.

We now turn to investigate the role of overage payments and adoption quality as triggers for switching. This analysis enables us to test Hypotheses 3 and 4. To do so, we utilize a proportional hazard regression in which the dependent variable, $h_{k,it}$, is customer i's switching hazard for event k (downward switch or upward switch), t time periods after adoption, given that the customer has not switched contracts by that time. In particular, we estimate the following duration model:

$$(2)\, h_{k,it} = h_{k0}(t) \cdot e^{\beta_k' Z_i(t)}, \quad k = (Upward\ switch, Downward\ switch); \quad i = 1 \dots n$$

where h_{k0} is the baseline hazard function for the event k (upward or downward switch), and $\boldsymbol{\beta}_k$ is the (event-specific) column vector of regression coefficients for event k. $Z_i(t)$ is a vector of covariates that may affect the hazard rate of individual i. $Z_i(t)$ includes the following set of variables:

$$(3)\, Z_i(t) = \left\{AdoptionQuality_{i,t-1}, Overage_{i,t-1}, \overline{Overage_{i,t-2}}, Activity_{i,t-1}, X_{i,t-1}\right\}$$

For each post-adoption period prior to switching, the main explanatory variables are customer *i's* adoption quality at time t-1, $AdoptionQuality_{i,t-1}$, and the amount paid in overage fees at time t-1, $Overage_{i,t-1}$. Adoption quality in this regression is a time varying variable that is calculated as the ratio between the monthly payment at time t and the mean pre-adoption payments. In addition, $\overline{Overage_{i,t-2}}$ represents the mean monthly amount paid prior to t-1.[9] We also control for the level of activity for the customer, $Activity_{i,t-1}$, and for other customer time-varying and non-time-varying characteristics in $X_{i,t\text{-}1}$.

We estimate the hazard models for upward and downward switching events across all adopting customers (including customers who did not eventually switch to a new contract after initial adoption). For each event we exclude from the estimation sample customers who underwent the other event type, and customers who eventually quit the new contract in favor of the old pay-per use scheme.[10]

The results, presented in Table 5, indicate that overage payments are positively associated with the hazard of upward switching and negatively associated with the hazard of downward switching. For the effect of adoption quality we find the opposite. Adoption quality at t-1 is negatively associated with upward switching at t, and positively associated with downward switching. These results suggest that customers who initially chose contracts that increased their monthly commissions (as compared with pre-adoption commissions) were more likely to conclude that they needed to switch to contracts with lower fixed payments, and were much less likely to switch to contracts with higher fixed payments. Thus, overall, our findings are consistent with the predictions of directional learning theory (Hypotheses 3 and 4).

In order to examine Hypothesis 2 we separately estimate the immediate influence of overage payments and their influence over longer periods of time. That is, we include in the hazard regression both the amount paid as overage at time t-1, and the mean monthly amount

[9] Although adoption quality exhibits high variation across customers, the variation of this variable over time is stable at the customer level. For instance, we find that the quality measure for 36% of the customers is identical over time. Furthermore, around 80% of downward-switchers have zero time variation in their quality measure. As a result, in the duration analysis we cannot split this measure into last period quality and previous periods' quality.

[10] We analyze quitting decisions in section IV.3

paid prior to t-1 (i.e., as of t-2). We see that the positive effect of overage payments on upward-switching decays quickly. Whereas one dollar paid in overage payments in month t-1 increases the hazard for upward-switching in month t by 8 percent, an increase of one dollar in the average amount paid as overage between adoption and t-2, increases the hazard of upward-switching by only 2 percent. This decrease in the hazard ratio indicates that if a specific overage payment does not lead to switching in the subsequent month it is much less likely to lead to switching in later months. These findings are generally consistent with Hypothesis 2.

The results presented thus far are not necessarily inconsistent with Bayesian learning theories. To examine whether standard Bayesian learning can explain the behavior we observe, we now turn to examining the consequences of switching (Hypothesis 1). In particular, we assume that choices that increase customers' monthly commissions (i.e., worsen their financial situations) are inconsistent with Bayesian learning.

V.2.2. Consequences of upward- and downward-switching

To evaluate the consequences of learning (Hypothesis 1), we exploit the longitudinal nature of our data and estimate the following panel data fixed effects regression:

$$(4)\quad Log(payment)_{it} = \beta_0 + \beta_D D_{Switching,it} + \beta_A Log(activity_{it} + 1) + \beta_X Log(X_{it} + 1) + \alpha_i + \eta_t + \varepsilon_{it}$$

where the variable, $D_{Switching,it}$ in Equation (4) is a dummy variable that equals one if customer i has switched contracts by time t and zero otherwise. Other variables are similar to the ones discussed above with regard to Equation (1). The regression results are reported in Table 6. In column 1 we report the estimation results for the entire sample of switching customers. The results indicate that, on average, customers' monthly commissions decreased by 3.2% after switching in comparison to post-adoption pre-switching payments. This result implies that learning indeed leads to better outcomes. However, this is not always the case. In columns (2) and (3) we split the sample of switching customers into downward- and upward-switchers, respectively. Our regression results indicate that customers who switched to contracts with smaller allowances were able to reduce their monthly payments by 30 percent. In contrast, the monthly payments of customers who switched to contracts with larger allowances increased by 11 percent after switching. Thus, while some switchers were able to substantially reduce their monthly commissions after switching, others systematically ended up paying more to the bank. This result contradicts a basic assumption in the Bayesian updating learning theory that learning cannot lead to worse economic outcomes (Hypothesis 1).

We find additional support for these patterns when we restrict attention to customers who switched contracts twice. We estimate Equation (4) for customers who switched contracts twice, considering only the months after the first switch.[11] The results, reported in Table 7, are qualitatively similar to the results presented in Table 6 for the first switch. Thus, as shown in column 1, customers who switched twice paid nearly 5 percent less after the second switch than they paid after the first switch. Nevertheless, while customers whose second switch was a downward switch reduced their monthly payments by 27.5 percent, customers whose second switch was an upward switch increased their monthly payments by 7.4 percent.

V.3. Quitting Decisions

We now turn to the analysis of quitting decisions and their implications, in order to further explore Hypothesis 5. Specifically, we investigate how overage payments and

[11] There are 127 customers who switched contracts twice. Among them 77 switched upward and 50 switched downward. Among the 77 second time upward-switchers, 61 customers paid overage payments between their first and second switches. In contrast, only 16 customers who paid overage payments between the two switches, switched downward.

experiences of poor adoption quality affect the tendency of adopting customers to quit the new subscription contracts in favor of the old pay-per-use scheme. We expect that both triggers will have a positive effect on quitting and, as before, that both these triggers will have a short-lived effect on the decision to quit. We estimate the following hazard regression:

$$(5)\, h_{q,it} = h_{q0}(t) \cdot e^{\beta_q' Z_i(t)}, \qquad i = 1 \dots n$$

The dependent variable in Equation (5) is customer i's quitting hazard t time periods after adoption, given that the customer has not quit by that time. h_{q0} is the baseline hazard function for quitting, and $\boldsymbol{\beta}_q$ is the column vector of regression coefficients for quitting. $Z_i(t)$ is similar to that in Equation (3).

Table 8 presents the regression results of the hazard regression for quitting. Higher overage payments in the month prior to quitting and lower adoption quality were positively associated with quitting. Interestingly, the mean level of overage payments up to two months before quitting had no significant effect on quitting. This result provides further support for the recency effect (Hypothesis 2), indicating that recent experiences have a strong influence on customers' decision making. It is difficult, however, to explain the observed behavior within the Bayesian learning framework. As noted above, customers who interpret overage payments as a signal for increased future consumption are not expected to quit three-part tariff contracts in favor of a pay-per-use scheme (Hypothesis 5).

V.4. Summary

Overall, our empirical analysis provides evidence that support Hypotheses 2, 3, and 4, yet contradicts Hypotheses 1 and 5.

We find that customers make decisions on the basis of recent information, and that if a trigger does not lead to switching in the subsequent month it is much less likely to lead to switching in later months (Hypothesis 2). We also find that low adoption quality and overage payments trigger contract switches (Hypothesis 3 and 4). At the same time, customers with low adoption quality are more likely to switch to contracts with lower fixed payments, and were much less likely to switch to contracts with higher fixed payments, whereas customers who experienced overage payments, were less likely to switch to contracts with lower fixed payments, and were much more likely to switch to contracts with higher fixed payments (Hypothesis 4). Finally, we find that some customers systematically ended up paying more to the bank after switching, contradicting Hypothesis 1, and that overage payments also trigger quitting decisions which cannot be explained within the Bayesian learning framework (Hypothesis 5).

VI. Discussion and Concluding Remarks

The assumption that individuals use all available relevant information when choosing among alternatives is fundamental for economic analysis. A common explanation for non-optimal choices is that deviations from the rational choice model are non-systematic and that market forces and experience should eventually correct for any inconsistencies in individual decision making. Accordingly, learning from experience is considered a primary vehicle through which individuals obtain relevant information and improve their choices. Yet alternative learning theories suggest that information processing is of a more subjective nature, stressing the role of personal experiences, especially recent ones, in shaping individuals' subsequent decisions. How, then, can we actually characterize individuals' learning?

In this paper, we analyzed the decision making processes of consumers faced with a menu of subscription contracts. Our empirical investigation examined both the initial adoption decision and possible subsequent contract changes. Using rich panel data spanning 30 months, including 6 months before the new contracts were introduced, we empirically investigated the determinants, the speed, and the consequences of learning. In general, our findings are most consistent with directional learning, and are inconsistent with Bayesian learning models. More specifically, we showed that learning is triggered by two distinct types of post-adoption experiences that involve clear reference points and are likely to be evaluated as negative by consumers. The first is an increase in monthly post-adoption payments as compared to pre-adoption payments (i.e., low adoption quality), and the second experience is payment of overage fees on top of the contract's fixed fee, in cases in which customers exceed their contracts' usage allowances. Customers dislike these unexpected experiences and switch contracts. Our findings indicate that, contrary to what Bayesian learning models would predict, customers' switching decisions may lead them to pay *higher* monthly commissions than they paid before switching.

This latter finding raises an essential question: If an effort to minimize monthly commission payments is not what drives the choice of a new contract, what then might explain contract choice in the case of switching? We postulate that customers switch to contracts that reduce the likelihood that the negative experience triggering the switch will recur. Downward switchers, on the one hand, aim to eliminate the experience of low adoption quality. Accordingly, following the switch, these customers reduce their payments by 30% compared with post-adoption pre-switching payments. This reduction is likely to cover the 17% increase in monthly commissions that these customers experience post-adoption (as compared with pre-adoption payments). Upward switchers, on the other hand, do not seem to focus on adoption quality. These customers' switching decisions are triggered by overage payments, and thus their new contract choices are focused on the elimination of this negative experience. We therefore expect upward switchers to choose contracts with fixed payments that (just) cover their overall payments prior to switching (including fixed and overage payments). Indeed, we find that the difference between the fixed payment of the new chosen contract and the payment in the month prior to switching is very close to zero ($0.17 on average). This difference increases as we move back in time from the month of switch ($0.22, $1.04, and $1.73, two, three, and four months prior to switching, respectively). Overall, these findings, together with our findings regarding the elapsed time to switch, are consistent with directional learning models, yet are difficult to explain within the standard Bayesian learning framework or reinforcement learning.

Our findings are particularly striking given the simple structure of the environment we study. Checking accounts cannot be regarded a new service for the customers we analyze. The average account tenure in our sample is 14 years. Moreover, and perhaps more importantly, the decisions we study do not require customers to consider the decisions of others. Accordingly, the failure of Bayesian updating to explain our findings does not hinge on the difficulty to form correct beliefs over other players' strategies.

Finally, our empirical findings could have direct bearing on the policy debate on whether regulatory intervention in subscription markets is warranted. In contrast to previous studies, our findings suggest that experienced customers might be worse off after gaining experience and hence that policy intervention is potentially warranted. Our findings about directional learning also highlight an underexplored aspect of the architecture of choice (Sunstein and Thaler 2008). While the debate around choice architecture typically concerns initial choices, the architecture of choice could very well also affect the consequences of learning that takes place after the initial choice. Naturally, further research is needed to better characterize the exact mechanisms through which regulators can improve on customers' decision making processes over time.

References

Agarwal, S., S. Chomsisengphet, C. Liu, N.S. Souleles. 2006. Do Consumers Choose the Right Credit Contracts? Working paper, The Wharton School, University of Pennsylvania.

Agarwal, S., J.C. Driscoll, X. Gabaix, D.I. Laibson. 2013. Learning in the Credit Card Market. Mimeo, Harvard University.

Ater, I. V. Landsman 2013. Do Customers Learn from Experience: Evidence from Retail Banking. *Management Science*, 59(9), 2019-2035.

Bordalo, P., N. Gennaioli, A. Shleifer. 2013. Salience and Consumer Choice. *Journal of Political Economy*, 121(5), 803-843.

Chiang, Yao-Min, David Hirshleifer, Yiming Qian, and Ann E. Sherman, 2011, Do Investors Learn from Experience? Evidence from Frequent IPO Investors, Review of Financial Studies 24, 1560–1589.

Choi, J., D. Laibson, B. Madrian, and A. Metrick. 2009. Reinforcement Learning and Savings Behavior. *Journal of Finance* 64:2515–34.

Crawford. V. 2013. Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games. *Journal of Economic Literature* 51:2, 512-527.

DellaVigna, S. 2009. Psychology and Economics: Evidence from the Field. *Journal of Economic Literature* 4**, 7**315–372.

DellaVigna, S., U. Malmendier. 2006. Paying Not to Go to the Gym. *American Economic Review*, 96(3), 694–719.

Erev, I., AE. Roth. 1998. Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria. *American Economic Review*, 88, 848-881.

Farber, H., and R. Gibbons. 1996. Learning and Wage Dynamics, *Quarterly Journal of Economics*, CXI, 1007-1047.

Gennaioli, N. and A. Shleifer. 2010. What Comes to Mind, *Quarterly Journal of Economics* 125(4): 1399-1433

Grether, David M. 1980: .Bayes Rule as a Descriptive Model: The Representative Heuristic. *Quarterly Journal of Economics*, November, 537-557.

Grubb, M.D., M. Osborne. 2014. Cellular Service Demand: Biased Beliefs, Learning and Bill-shock. *American Economic Review* (forthcoming).

Harstad, R.M., R. Selten. 2013. Bounded-Rationality Models: Tasks to Become Intellectually Competitive. *Journal of Economic Literature* **51**:2, 496-511.

Haselhuhn, M.P., D.G. Pope, M.E. Schweitzer, P. Fishman. 2012. The Impact of Personal Experience on Behavior: Evidence from Video-rental Fines. *Management Science*, 58(1), 52–61.

Hertwig, R., G. Barron, E.U. Weber, and I. Erev. 2004. Decisions from Experience and the Effect of Rare Events in Risky Choice," *Psychological Science*,15, 534–539.

Herweg, F., K. Mierendorff. 2013. Uncertain Demand, Consumer Loss Aversion, and Flat-Rate Tariffs. *Journal of the European Economic Association*, 11(2), 399-432.

Hogarth, R.M., H.J. Einhorn. 1992. Order Effects in Belief Updating: The belief-adjustment Model. Cognitive Psychology, 24, 1–55.

Ito, K. 2014. Do Consumers Respond to Marginal or Average price? Evidence from Nonlinear Electricity Pricing. *American Economic Review* 104(2), 537–563.

Kahneman, D., A. Tversky. 1972. Subjective Probability: A Judgment of Representativeness. *Cognitive Psychology*, 3, 430–454.

Kaustia, Markku and Samuli Knüpfer, 2008, Do Investors Overweight Personal Experience? Evidence from IPO subscriptions, Journal of Finance 63, 2679-2702.
Ketcham, J., C. Lucarelli, E.J. Miravete, M.C. Roebuck. 2012. Sinking, Swimming, or Learning to Swim in Medicare Part D. *American Economic Review*, 102(6) 2639–2673.
Lambrecht, A., K. Seim, B. Skiera. 2007. Does uncertainty matter? Customer behavior under three-part tariffs. *Marketing Science* **26**(5) 698–710.
Landsman, V., M. Givon. 2010. The diffusion of a new service: Combining Service Consideration and Brand Choice. *Quantitative Marketing and Economics* **8**(1) 91–121.
Levitt, S.D., J.A. List. 2008. Homo Economicus Evolves. *Science*, 319(5865), 909–10.
List, J.A. 2003. Does Market Experience Eliminate Market Anomalies? *Quarterly Journal of Economics*, 118(1), 41–71.
List, J.A. 2004. Neoclassical Theory versus Prospect Theory: Evidence from the Marketplace. *Econometrica*, 72(2), 615–25.
Malmendier, U., S. Nagel. 2011, Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking? Quarterly Journal of Economics, 126, 373-414
Miravete, E.J. 2003. Choosing the Wrong Calling Plan? Ignorance and Learning. *American Economic Review*, 93(1), 297–310.
Ockenfels, A., R. Selten. 2014. Impulse balance in the newsvendor game, *Games and Economic Behavior*, 86, July, 237–247.
Ortoleva, Pietro. 2012. Modeling the Change of Paradigm: Non-Bayesian Reaction to Unexpected News, American Economic Review 102, 2410—2436.
Rabin, M. 2000, Risk aversion and expected-utility theory: A calibration theorem, Econometrica 68(5), 1281-1292.
Rabin, Matthew. 2002. Inference by Believers in the Law of Small Numbers. *Quarterly Journal of Economics* 117 (3): 775–816.
Roth, A.E. and I. Erev 1995. Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. Games and Economic Behavior, 8(1), pp. 164-212.
Selten, R., J, Buchta. 1999. Experimental sealed bid first price auctions with directly observed bid functions. In: Budescu, Erev, Zwick (Eds.), Games and Human Behavior: Essays in the Honor of Amnon Rapoport. Lawrenz Associates, Mahwah, NJ, pp.79–104.
Selten, R., T. Abbink, and Cox, R. 2005. Learning direction theory and the winner's curse. *Experimental Economics* 8(1),5–20.
Selten, R., T. Chmura, 2008. Stationary concepts for experimental 2×2-games. *American Economic Review* .98(3), 938–966. Selten, R., and R. Stoecker. 1986.: End Behavior in Sequences of Finite Prisoners' Dilemma Supergames: A Learning Theory Approach. *Journal of Economic Behavior and Organization*, 7, 47-70
Smith, M. J. 1982: Evolution and the Theory of Games. Cambridge, United Kingdom: Cambridge University Press.
Sobel, J. 2000. Economists' Models of Learning, Journal of Economic Theory 94(1): 241-261.
Sunstein, C., R. Thaler. 2008. *Nudge: Improving Decisions About Health, Wealth, and Happiness.* New Haven: Yale University Press.
Thorndike, E. L. 1932. *The Fundamentals of learning*, New York: AMS Press.
Tversky, A., D. Kahneman. 1973. Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology*, 5(2), 207-32.

Woodford, M. 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton: Princeton University Press.

TABLE 1 – EXAMPLE OF THREE-PART TARIFF CONTRACTS

Contract #	*Monthly payment*	*Overage payment*		*Allowance*		
		Clerk-assisted activities	*Direct activities/ Check deposits*	*Clerk-assisted activities*	*Direct activities*	*Check deposits*
1	$4.75	$1.2	$0.30	0	7	7
5	$9.50	$1.2	$0.30	7	Unlimited	12

TABLE 2A – CHOICE OPTIMALITY EX-ANTE FOR CONTRACT ADOPTERS

Optimal / ***Chosen***	*Contract 1*	*Contract 2*	*Contract 3*	*Contract 4*	*Contract 5*	*Contract 6*	
Contract 1	**92.8%**	0.3%	5.9%	0.8%	0.2%	0.0%	100.0%
Contract 2	94.5%	**3.2%**	0.0%	2.1%	0.2%	0.0%	100.0%
Contract 3	74.8%	0.0%	**23.9%**	0.0%	1.2%	0.0%	100.0%
Contract 4	89.2%	2.4%	0.0%	**7.5%**	0.8%	0.0%	100.0%
Contract 5	75.6%	4.6%	3.1%	12.0%	**4.5%**	0.2%	100.0%
Contract 6	69.1%	7.0%	1.2%	12.6%	7.9%	**2.1%**	100.0%

TABLE 2B – CHOICE OPTIMALITY EX-POST FOR CONTRACT ADOPTERS

Optimal / ***Chosen***	*Contract 1*	*Contract 2*	*Contract 3*	*Contract 4*	*Contract 5*	*Contract 6*	
Contract 1	**94.8%**	0.3%	4.0%	0.7%	0.2%	0.0%	100.0%
Contract 2	93.9%	**4.0%**	0.0%	2.0%	0.1%	0.0%	100.0%
Contract 3	80.4%	0.0%	**18.7%**	0.0%	0.8%	0.0%	100.0%
Contract 4	89.6%	3.1%	0.0%	**6.4%**	0.9%	0.0%	100.0%
Contract 5	77.3%	5.2%	3.0%	10.9%	**3.5%**	0.2%	100.0%
Contract 6	70.8%	7.7%	1.1%	11.6%	6.4%	**2.4%**	100.0%

The numbers represent the distribution of optimal contracts across adopters for each of the available contracts. Each row presents the percent distribution of optimal contracts for the customers who chose the particular contract represented in that row. Optimality is calculated based on 3 months after adoption. When we aggregate over all adopting customers we find that 29% chose their cost-minimizing contracts. Further details are provided in online Appendix A.

TABLE 3 – DESCRIPTIVE STATISTICS: ADOPTERS, SWITCHERS AND QUITTERS

Variable	*(1)* *Adopters* Mean (St. Dev.)	*(2)* *Downward-switchers* Mean (St. Dev.)	*(3)* *Upward-switchers* Mean (St. Dev.)	*(4)* *Quitters* Mean (St. Dev.)
Pre-adoption mean payments [a]	7.54 (3.45)	8.52 (3.23)	7.62 (3.32)	6.59 (3.17)
Ratio of mean monthly payments (after and before adoption)	1.17 (0.45)	1.30 (0.63)	1.12 (0.44)	1.34 (0.52)
Mean overage payments (after adoption and before switch)	0.43 (1.17)	0.21 (0.74)	1.81 (2.66)	0.62 (1.65)
Mean time to switch/quit (months)	n/a	5.68 (4.41)	8.9 (5.64)	
Account tenure (years)	13.78 (9.32)	13.85 (9.25)	13.00 (9.16)	13.62 (9.58)
Age of youngest account holder	44.41 (13.71)	43.34 (12.66)	44.28 (14.29)	46.55 (15.39)
Number of account owners	1.44 (0.51)	1.52 (0.50)	1.40 (0.51)	1.40 (0.52)
Parental Social Security benefits (for children below age of 18) in thousand U.S. dollars [b]	0.04 (0.11)	0.05 (0.13)	0.05 (0.13)	0.03 (0.09)
Elderly Social Security benefits in thousands of US dollars [b]	0.07 (0.21)	0.06 (0.19)	0.08 (0.24)	0.11 (0.25)
Number of salaries [a]	0.74 (0.79)	0.82 (0.81)	0.64 (0.77)	0.59 (0.72)
Socio-economic measure of residence of account holder (scale of 1–10) [b]	5.17 (2.23)	5.29 (2.14)	5.05 (2.20)	5.17 (2.25)
Mean number of account information inquiries [b]	7.15 (14.08)	8.68 (13.77)	7.01 (12.13)	5.64 (10.68)
Mean number of clerk-assisted transactions [b]	0.95 (1.23)	1.06 (1.18)	1.24 (1.41)	1.11 (1.29)
Mean number of transactions through direct channels [b]	3.35 (3.68)	3.74 (3.72)	2.94 (3.47)	2.31 (2.97)
Mean number of check transactions [b]	3.59 (5.08)	4.51 (4.85)	3.86 (5.40)	3.02 (5.08)
Mean number of marketing calls [b]	0.08 (0.08)	0.08 0.07	0.08 (0.09)	0.09 (0.09)
Customers	32650	827	1441	2068

[a] Calculated based on all the months before adoption.
[b] Calculated based on three months before adoption.
Due to confidentiality concerns we are not allowed to reveal the summary statistics for the salary, loans and savings variables. We use these variables in the regression analysis.

TABLE 4 – CUSTOMER ADOPTION PAYMENT REGRESSIONS

	(1)	(2)	(3)	(4)
VARIABLES	***All customers***	***All switchers***	***Upward switchers***	***Downward switchers***
Post adoption month	0.089*** (0.002)	0.054*** (0.007)	-0.001 (0.009)	0.169*** (0.013)
Number of clerk-assisted transactions	0.128*** (0.001)	0.156*** (0.004)	0.195*** (0.005)	0.097*** (0.006)
Number of direct transactions	0.037*** (0.001)	0.037*** (0.003)	0.040*** (0.004)	0.039*** (0.006)
Number of check deposits	0.064*** (0.001)	0.078*** (0.004)	0.088*** (0.005)	0.070*** (0.006)
Number of owners	0.054*** (.011)	0.068 (0.043)	0.080 (0.048)	0.028 (0.073)
Parental Social Security benefits	0.345*** (0.023)	0.299*** (0.065)	0.188** (0.076)	0.554*** (0.124)
Elderly Social Security benefits	0.045*** (0.013)	-0.034 (0.048)	-0.016 (0.057)	-0.061 (0.066)
Number of information inquiries	0.001 (0.001)	0.009*** (0.003)	0.007** (0.003)	0.009** (0.004)
Salary	-0.001 (0.002)	0.009 (0.011)	0.005 (0.014)	0.002 (0.017)
Number of salaries	-0.010*** (0.003)	-0.031** (0.013)	-0.020 (0.016)	-0.033 (0.019)
Loans	0.023*** (0.001)	0.031*** (0.007)	0.041** (0.009)	0.022** (0.009)
Savings	0.012*** (0.001)	0.022*** (0.005)	0.029*** (0.006)	0.015** (0.007)
Constant	1.840*** (0.005)	1.814*** (0.018)	1.733*** (0.019)	1.930*** (0.033)
Observations	923673	46767	29,847	16,920
R-squared	0.357	0.367	0.406	0.336
Number of customers	32650	2268	1,441	827

*** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses are clustered at the account level.
The dependent variable in all regressions is the (log) monthly payment to the bank. An observation is an account/month. All regressions include individual account and month fixed effects. The estimation results include only the months before contract switching (including pre-adoption months, Equation (1) in the text). The sample of customers in column 1 includes all customers. In column 2 we focus on switching customers and in columns 3 and 4 on upward and downward switchers, respectively.

TABLE 5 – HAZARD REGRESSION ANALYSIS FOR SWITCHING DECISION

Variable	*Downward-Switching Hazard Regression* Parameter Estimate (Standard Error)	Hazard ratio	*Upward-Switching Hazard Regression* Parameter Estimate (Standard Error)	Hazard ratio
***Adoption Quality*[a]**	0.01*** (0.00)	1.01	-0.12*** (0.04)	0.89
***Overage*[a]**	-0.09*** (0.02)	0.91	0.08*** (0.00)	1.08
***Mean Past Overage*[a]**	-0.07*** (0.02)	0.94	0.02*** (0.00)	1.02
***Loans*[a]**	0.00 (0.01)	1.00	0.00 (0.00)	1.00
***Parental Social Security benefits (for children below the age of 18)*[a]**	1.17*** (0.36)	3.23	0.88*** (0.28)	2.41
***Elderly Social Security benefits*[a]**	-0.08 (0.20)	0.92	0.24*** (0.12)	1.28
***Monthly number of clerk-assisted transactions*[a]**	0.00 (0.03)	1.00	-0.10*** (0.02)	0.91
***Monthly number of direct transactions*[a]**	-0.02** (0.01)	0.98	0.01 (0.01)	1.01
***Monthly number of check transactions*[a]**	0.04*** (0.01)	1.04	-0.08*** (0.00)	0.93
***Number of owners*[a]**	0.38*** (0.07)	1.47	0.14*** (0.05)	1.15
***Number of salaries*[a]**	0.01 (0.05)	1.01	-0.02 (0.04)	0.98
***Savings*[a]**	0.00 (0.00)	1.00	0.00 (0.00)	1.00
***Salary*[a]**	-0.08 (0.06)	0.93	-0.10*** (0.04)	0.90
***Monthly number of account status inquiries*[a]**	0.00 (0.00)	1.00	0.00 (0.00)	1.00
Account tenure	0.00 (0.00)	1.00	-0.01 (0.00)	0.99
Socio-economic indicator(scale of 1 to 10)	0.02 (0.02)	1.02	-0.02 (0.01)	0.98
Pre adoption number of marketing calls	0.10 (0.02)	1.11	0.02** (0.02)	1.02
Customer age	0.00 (0.00)	1.00	0.00 (0.00)	1.00
***Number of events*[c]**	827		1,441	

[a]Time-varying variable

[b] Non-time-varying variable

[c] customers who underwent the other event type (up-switch for the down switching regression and down-switching for the up switching regression), and customers who eventually quit the contracts were excluded from the respective sample

*** p<0.01, ** p<0.05, * p<0.1

The table shows the results of the duration analysis (Equation 2) for the upward and downward switching decisions. The hazard ratio corresponds to the exponentiated coefficient. That is, if *b* is greater (smaller) than one, the difference (*b*−1)*100 indicates the percentage by which a one-unit increase in the explanatory variable would increase (decrease) the hazard of a switch. As is standard in survival analyses, we also present the original coefficients and the standard errors for the original coefficients.

TABLE 6 – CUSTOMER SWITCHING PAYMENT REGRESSIONS

VARIABLES	***All switchers*** *(over post-adoption months)*	***Upward switchers*** *(over post-adoption months)*	***Downward switchers*** *(over post-adoption months)*
Post switching month	-0.030***	0.119***	-0.325***
	(0.006)	(0.004)	(0.008)
Number of clerk-assisted transactions	0.138***	0.152***	0.093***
	(0.004)	(0.004)	(0.006)
Number of direct transactions	0.009***	0.010***	0.005
	(0.002)	(0.002)	(0.003)
Number of check deposits	0.055***	0.056***	0.044***
	(0.004)	(0.004)	(0.009)
Number of owners	0.073**	0.039	0.135**
	(0.035)	(0.036)	(0.059)
Parental Social Security benefits	0.023	0.006	-0.203*
	(0.053)	(0.047)	(0.108)
Elderly Social Security benefits	0.088**	0.035	0.023
	(0.038)	(0.026)	(0.036)
Number of information inquiries	0.007***	0.006***	0.010***
	(0.002)	(0.002)	(0.002)
Salary	0.008	0.003	-0.003
	(0.008)	(0.007)	(0.013)
Number of salaries	0.012	0.015*	0.012
	(0.009)	(0.008)	(0.011)
Loans	0.025***	0.004	0.018***
	(0.005)	(0.004)	(0.007)
Savings	0.015***	0.015***	0.003
	(0.003)	(0.003)	(0.005)
Constant	1.863***	1.848***	2.084***
	(0.019)	(0.020)	(0.037)
Observations	36,006	24,379	11,627
R-squared	0.273	0.421	0.364
Number of customers	2,268	1,441	827

*** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses are clustered at the account level.
The dependent variable in all regressions is the (log) monthly payment to the bank. An observation is an account/month.
All regressions include individual account and month fixed effects. The estimation results include only the months after contract adoption (Equation (4) in the text).

TABLE 7 –PAYMENT REGRESSIONS FOR SECOND-TIME SWITCHERS

	(1)	(2)	(3)
VARIABLES	All ‘second-time’ switchers	‘Second-time’ upward switchers	‘Second-time’ downward switchers
Post switching month	-0.047*	0.074***	-0.275***
	(0.027)	(0.023)	(0.041)
Number of clerk-assisted transactions	0.100***	0.091***	0.114***
	(0.020)	(0.019)	(0.037)
Number of direct transactions	0.011	0.024	0.005
	(0.013)	(0.017)	(0.004)
Number of check deposits	0.068***	0.092***	0.045***
	(0.018)	(0.026)	(0.010)
Number of owners	0.062	-0.024	0.156**
	(0.100)	(0.082)	(0.063)
Parental Social Security benefits	-0.445**	-0.117	-0.194*
	(0.215)	(0.162)	(0.108)
Elderly Social Security benefits	0.489**	-0.321	0.021
	(0.026)	(0.981)	(0.037)
Number of information inquiries	0.006	0.012	0.009***
	(0.012)	(0.012)	(0.003)
Salary	-0.023	0.008	-0.008
	(0.023)	(0.022)	(0.013)
Number of salaries	0.067*	0.036	0.012
	(0.010)	(0.035)	(0.014)
Loans	0.008	-0.012	0.022***
	(0.017)	(0.018)	(0.008)
Savings	0.033**	0.021	0.003
	(0.013)	(0.013)	(0.006)
Constant	1.880***	1.791***	2.079***
	(0.125)	(0.089)	(0.037)
Observations	1572	987	585
R-squared	0.136	0.223	0.353
Number of customers	127	77	50

*** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses are clustered at the account level.
The dependent variable in all regressions is the (log) monthly payment to the bank. An observation is an account/month. All regressions include individual account and month fixed effects. The estimation results include only the months after the first switching decision. The sample of customers shown in column 1 includes all the customers who switched twice. In columns 2 we focus on customers whose second switch was to a contract with larger allowance, while in column 3 we focus on customers whose second switch was to a contract with a smaller allowance.

TABLE 8 – HAZARD REGRESSION ANALYSIS FOR QUITTING DECISION

	Parameter Estimate (Standard Error)	***Hazard ratio***
Adoption Quality[a]	0.02***	1.02
	(0.00)	
Overage[a]	0.06***	1.06
	(0.00)	
Mean Past Overage[a]	0.01	1.01
	(0.01)	
Loans[a]	-0.03***	0.97
	(0.01)	
Parental Social Security benefits (for children below the age of 18)[a]	-1.63***	0.20
	(0.40)	
Elderly Social Security benefits[a]	-0.01	0.99
	(0.11)	
Monthly number of clerk-assisted transactions[a]	-0.12***	0.89
	(0.02)	
Monthly number of direct transactions[a]	-0.13***	0.88
	(0.01)	
Monthly number of check transactions[a]	-0.06***	0.95
	(0.01)	
Number of owners[a]	-0.07	0.93
	(0.05)	
Number of salaries[a]	-0.18***	0.83
	(0.04)	
Savings[a]	0.00***	1.00
	(0.00)	
Salary[a]	-0.11***	0.89
	(0.04)	
Monthly number of account status inquiries[a]	-5.83E-04	1.00
	(0.00)	
Account tenure	-2.13E-03	1.00
	(0.00)	
Socio-economic indicator(scale of 1 to 10)	0.02	1.02
	(0.01)	
Pre adoption number of marketing calls	0.08***	1.08
	(0.02)	
Customer age	3.90E-03**	1.00
	(0.00)	
Number of events	2,068	

[a]Time-varying variable

[b] Non-time-varying variable

*** p<0.01, ** p<0.05, * p<0.1

The Table shows the results of the duration analysis (Eq. 5) for the quitting decisions. The hazard ratio corresponds to the exponentiated coefficient. That is, if *b* is greater (smaller) than one, the difference (b−1)*100 indicates the percentage by which a one-unit increase in the explanatory variable would increase (decrease) the hazard of quitting. As is standard in survival analyses, we also present the original coefficients and the standard errors for the original coefficients.

FIGURE 1A – ADOPTION QUALITY AMONG ALL ADOPTERS

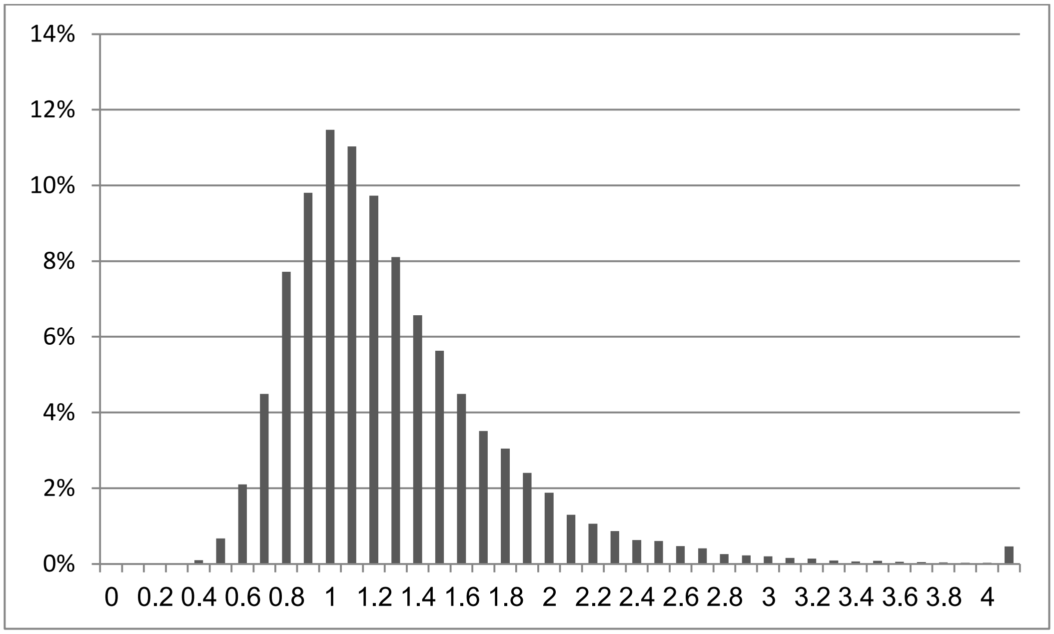

FIGURE 1B – ADOPTION QUALITY AMONG UPWARD- AND DOWNWARD-SWITCHERS

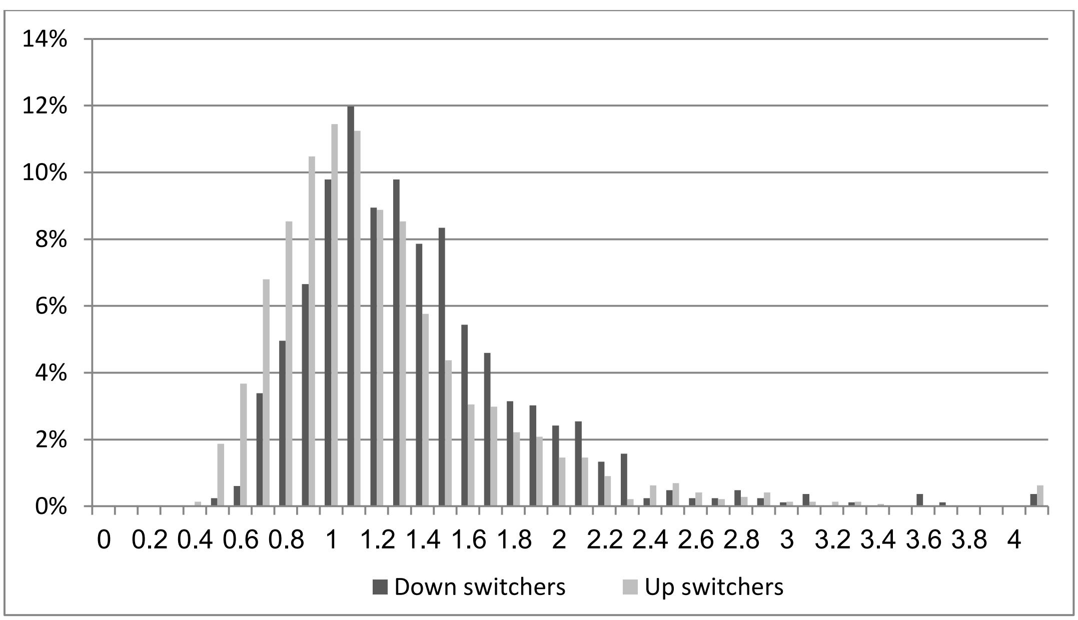

Speed of Switching Decisions

FIGURE 2 – CUMULATIVE DISTRIBUTION OF TIME UNTIL SWITCH

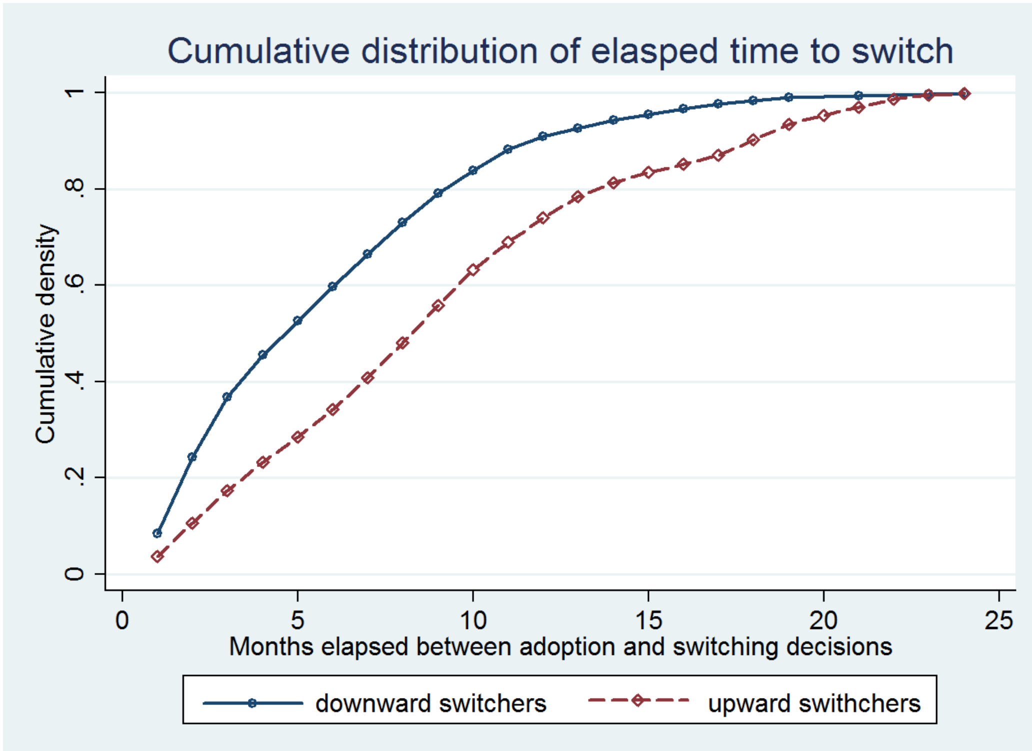

FIGURE 3 – HISTOGRAM OF TIME BETWEEN LAST OVERAGE PAYMENT AND SWITCH

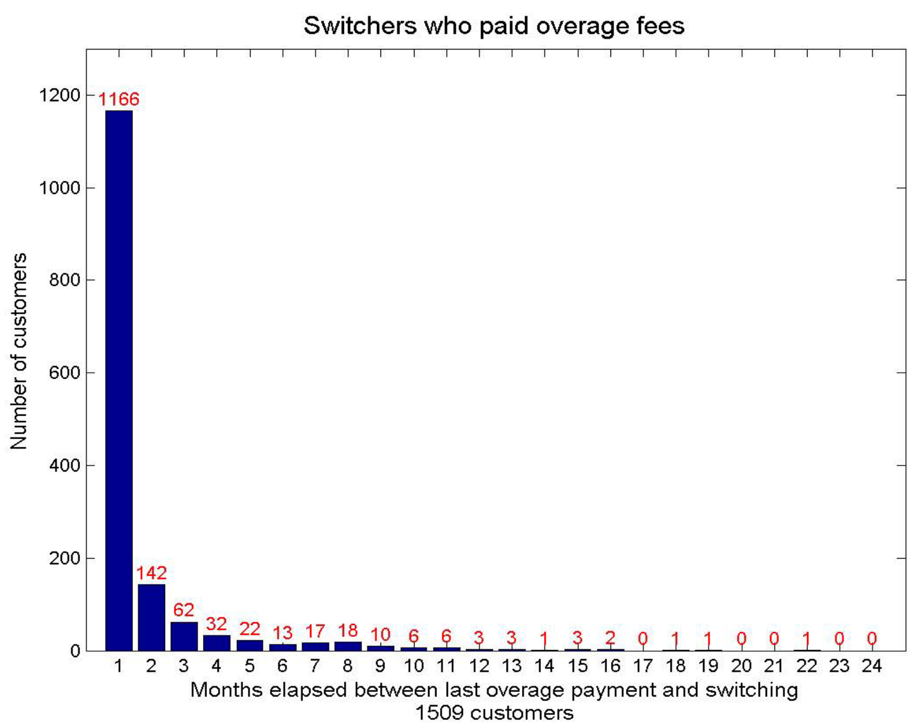

Speed of Quitting Decisions

FIGURE 4 – CUMULATIVE DISTRIBUTION OF TIME UNTIL QUITTING

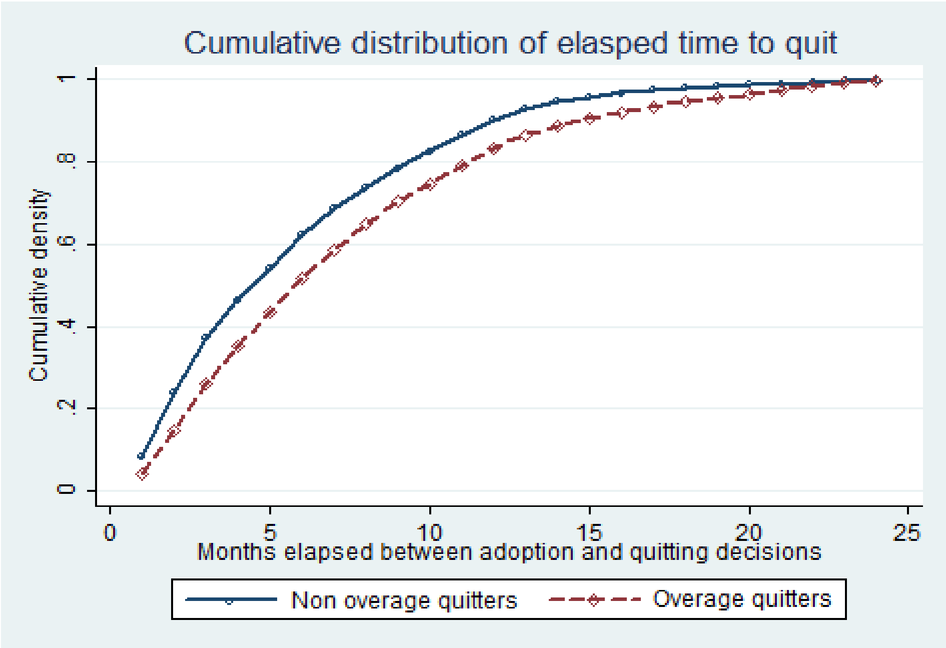

Online Appendix A – Optimality Calculation Schemes

A customer might calculate the cost of a plan in a given three-part tariff plan menu in several ways. For example, customers might consider only their activity in the month of choice while evaluating the plans. Alternatively, they might take into account a longer period of time spanning several months of activity. Moreover, customers might compare the overall payments over the entire time period across all plans based on their activity in each month, or alternatively calculate the payment associated with each plan according to their mean monthly activity levels. Because we are not aware of the actual methods used by customers to calculate and compare payments across different plans, we employed several payment calculation approaches to evaluate choice optimality. First, we considered different time periods for the optimality calculation, ranging from one month to six months. Second, we used two calculation schemes to calculate the payment associated with each plan. The first scheme was based on the monthly mean number of transactions (according to each of the transaction types) over the relevant timeframe. The second was based on the overall payment for each plan based on the customer's actual usage over the relevant timeframe. While the latter approach is more accurate in terms of plan optimality, it is more complex to compute. Take, for example, a customer who uses the account heavily only once a year. This customer might do best by choosing a cheap plan with a small allowance and just paying the overage payments during the month of heavy usage. But if that same customer calculates her average monthly activity (taking that month into account), she might conclude that she needs a larger allowance, and she will end up buying a more expensive plan and paying larger amounts each month. Third, a customer might choose a plan that is not optimal for his or her past usage behavior, and yet can be optimal given a behavioral change. We, therefore, assessed the optimality of customers' plan choices using both an ex-ante approach (i.e., by evaluating pre-adoption usage behavior) and an ex-post approach (i.e., evaluating post-adoption usage behavior). The ex-post criterion might be a more accurate criterion for assessing optimality if, at the time of adoption, customers take into account their expected changes in usage behavior. We find that our optimality assessments are rather similar under the different schemes. Therefore, in the paper we present our analysis results based on evaluation of the monthly mean number of transactions over three months (i.e., not the overall payment), either ex-ante or ex-post, depending on the analysis.

In Table A1 we present 24 different optimality calculation schemes used in our plan choice optimality assessment. The optimality calculation schemes differ on three levels: (1) the length of the timeframe investigated in order to assess optimality, (2) the basis for optimality calculation (i.e., overall payment or mean monthly activity level), and (3) the ex-post or ex-ante assessment. Table A1 presents the 24 different calculation schemes.

TABLE A1 – OPTIMALITY CALCULATION SCHEMES

	Reference time period	***Length of time period***	***Calculation basis***
1.	ex-ante	1 month[a]	overall payment
2.	ex-ante	1 month[a]	mean activity level
3.	ex-ante	2 months	overall payment
4.	ex-ante	2 months	mean activity level
5.	ex-ante	3 months	overall payment
6.	ex-ante	3 months	mean activity level
7.	ex-ante	4 months	overall payment
8.	ex-ante	4 months	mean activity level
9.	ex-ante	5 months	overall payment
10.	ex-ante	5 months	mean activity level
11.	ex-ante	6 months	overall payment
12.	ex-ante	6 months	mean activity level
13.	ex-post	1 month[a]	overall payment
14.	ex-post	1 month[a]	mean activity level
15.	ex-post	2 months	overall payment
16.	ex-post	2 months	mean activity level
17.	ex-post	3 months	overall payment
18.	ex-post	3 months	mean activity level
19.	ex-post	4 months	overall payment
20.	ex-post	4 months	mean activity level
21.	ex-post	5 months	overall payment
22.	ex-post	5 months	mean activity level
23.	ex-post	6 months	overall payment
24.	ex-post	6 months	mean activity level

[a] For optimality assessments based on a 1-month period there is no difference in calculated plan payments between the two calculation bases (overall payment and mean activity level).

Next, we provide an example to illustrate the difference between the optimality calculation schemes. Take, for example, a customer who performed the following numbers of transactions in each of the three transaction types over three months (Table A2):

TABLE A2 – ACTIVITY DESCRIPTION

	Clerk-assisted transactions	*Direct transactions*	*Check transactions*
Month 1	0	7	10
Month 2	17	8	8
Month 3	1	5	11
Mean activity	6	6.66	9.66

Table A3 presents the calculation of the plan payment using the overall payment calculation basis for the four plans that were available in those three months, and for the old pay-per-use payment system. According to Table A3, Plan 1 is the optimal plan because the overall payment for this plan is the lowest.

TABLE A3 – PAYMENT CALCULATION USING OVERALL PAYMENT AS THE CALCULATION BASIS

	Plan 1	Plan 2	Plan 3	Plan 4	old system
Month: 1	$5.65	$6.25	$7.65	$14.00	$13.51
Month: 2	$26.60	$27.50	$24.55	$14.00	$ 6.16
Month: 3	$7.20	$7.80	$7.95	$14.00	$4.78
Sum	**$39.45**	$41.55	$40.15	$42.00	$24.45

Table A4 presents the calculation of the plan payment using the mean monthly activity level calculation scheme. According to this calculation scheme the optimal plan is Plan 3.

TABLE A4 – PAYMENT CALCULATION USING MEAN ACTIVITY LEVEL AS THE CALCULATION BASIS

	Plan 1	Plan 2	Plan 3	Plan 4	No plan
Payment according to mean activity level	$13.05	$13.75	**$11.30**	$11.50	$13.51