Testing Alternative Learning Theories: Evidence from Subscription Contracts

Itai Ater — Tel Aviv University

Vardit Landsman — Tel Aviv University and Erasmus University

March 27, 2016

Abstract

Analyzing panel data of 32,650 checking-account holders facing a menu of three-part tariff contracts, we document several findings that indicate that subscribers use simple heuristics to learn about the desirability of the contracts they have chosen. Our main findings are: subscribers change contracts in a direction that diminishes the probability of re-experiencing the trigger for switching; subscribers exhibit recency effects in switching, and after switching the majority of switchers systematically pay higher fees than they did before. We argue that directional learning theory could explain why consumers behave in a manner that yields suboptimal economic outcomes.

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

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

* Itai Ater is a Senior Lecturer of Economics and Strategy at the School of Management at Tel-Aviv University, Israel, Ater@post.tau.ac.il. Vardit Landsman is a Senior Lecturer of Marketing at the School 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 and seminar participants at Tel Aviv University, Hebrew University, Bar Ilan University and the Interdisciplinary Center (Herzliya). We also thank Emanuel Marcu for excellent research assistance. Remaining errors are our own.

1. 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 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? Does Bayesian learning offer a reasonable approximation to the observed behavior in a real-market environment? Which learning theory best describes behavior?

Empirically distinguishing among learning theories is challenging, for several reasons. First, in order to identify learning mechanisms that are at play in repeated decision situations, it is necessary to obtain data that 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 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 an individual's choice against its alternatives. Third, different learning theories often provide similar predictions, such that it is difficult to use empirical results 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 empirical studies have investigated which learning theory best explains individual learning behavior.

We address this gap in the literature in a study that, to the best of our knowledge, is the first to use field data to test the extent to which specific learning theories, including directional learning theory, explain the dynamic behavior of individuals. Specifically, we take 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 many customers do not choose their optimal contracts at the time of subscription selection. We study decisions made by a group of 32,650 holders of checking accounts in a large retail bank, 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 subsequent switching decisions (or lack thereof). Overall, there are 2,268 switchers in our data, and 2,030 customers who returned to the old pay-per-use pricing scheme. The data also include rich information on each account holder, including age, social security payments, salary, loans, savings and account tenure.

We document the following empirical findings. First, we show that 69% of the customers adopted contracts with excessively large allowances compared to their cost-minimizing contracts. This pattern does not depend on whether our calculations are based on customers' pre- or post- contract-adoption usage. Second, we identify two factors that trigger contract switching: *overage payments*, incurred when a customer exceeds his or her contract allowance, and high *post-pre commission ratio*, measured as the ratio between post- and pre-adoption commission payments. Thus, customers who on a given month paid more than their monthly fixed payments, or paid more than they had paid before adopting a three-part tariff contract were more likely to switch contracts. Third, we find that the switching decision takes place shortly after the respective trigger is realized, and that the impact of these triggers decays quickly, an observation that is consistent with the recency effect (Hogarth and Einhorn 1992). Fourth, we observe that the two triggers we identify have opposite effects on the direction of the subsequent contract switch. Overage payments are associated with an upward switch (a switch to a contract with higher fixed fees and larger allowances), whereas a high post–pre commission ratio is associated with a downward switch (i.e., a switch to a contract with lower fixed fees and smaller allowances). A fifth related finding is that these directional switching decisions result in opposite effects on the customer's average monthly payment to the bank. Customers who switched upward paid higher commission payments to the bank after the switch, whereas those who switched downward paid less after they switched. This finding is especially important because, rationally, one would expect a customer to switch only to a contract that lowers his or her average monthly payments. Sixth, we do not find indications for learning among consumers who do not switch contracts. That is, customers who do not switch contracts do not adjust their consumption behavior in order to more fully utilize their allowance of 'free' transactions. Specifically, we do not see an evolution of usage behavior in which the number of transactions becomes concentrated at allowance limits over time. Finally, we find that the likelihood of quitting, i.e., abandoning the three-part tariff contract and returning to the original pay-per-use scheme, is also positively associated with deviations from the monthly fixed payments (i.e., overage payments) and payment increases relative to pre-adoption payments.

Our conjecture based on these findings is that customers use simple heuristics to learn about the desirability of their contract choices. In particular, we propose that customers behave according to the following decision rule: the monthly bill is compared with a reference payment. If the bill is greater than the reference payment, the customer quickly adjusts the contract choice in a way that diminishes the likelihood of once again exceeding the reference payment. We propose that this heuristic is a simplified version of the behavior implied by directional learning theory. To substantiate this claim, in Section 2 we describe main features of Bayesian, reinforcement and directional learning theories. We then develop testable hypotheses implied by these theories, and in doing so we emphasize instances in which the predictions of the theories clash. In the data analysis presented in Section 3 we test these hypotheses, focusing on customers' decisions to update their initial contract choices, which we consider as an indication of learning.

As noted above, we find that the majority of switching customers paid higher average monthly payments after they switched; that is, they made the "wrong" switching decisions. A common approach to explain deviations from classic rational choice behavior is the *biased belief* approach, which relaxes the assumption of individuals' rational expectations regarding their future usage of a given product or its expected quality (Grubb 2015a, 2015b). Busse et al. (2015), for example, explore the impact of weather on purchasing decisions, and find evidence suggesting that projection bias induces individuals to purchase convertibles on warm, sunny days and to purchase 4-wheel-drive vehicles on cold-weather days, in a way that is inconsistent with classical utility theory. In a paper that is more related to the setting of the current study, Grubb and Osborne (2015), investigate contract choices among new cellular subscribers. They present evidence that cost-minimizing subscribers are overconfident and underestimate the variance of future calling. The biased belief approach seems rather natural in contexts in which consumers face high uncertainty, such as choices among new products or services. In our data set, however, the average tenure of customers we investigate is 14 years. Therefore, biased beliefs regarding future usage are probably not likely to have a major role in explaining customers' behavior in our setting.

Another common approach to explain deviations from the standard model is based on the view that individuals have non-standard preferences, and hence these deviations are actually not "wrong". Ater and Landsman (2013), for example, use a data set similar to the one used herein—namely, data on consumer checking accounts—and assume that learning must move individuals in the "right" direction. They therefore infer that consumers are not trying to minimize costs, but rather exhibit *non-standard preferences*, being *loss-averse* or *overage-averse*. While the approach of Ater and Landsman (2013) is plausible, it is also subject to the critique that it is unlikely that, in the specific context of managing a checking account, consumers have objectives, or preferences, other than cost minimization. Additionally, the coefficient of loss aversion in Ater and Landsman (2013) which is equal to 3.5 is greater than the conventional measure of loss aversion found in other studies (i.e. 2.25 as documented in Tversky and Kahneman 1992 and other studies). This also raises doubts on whether non-standard preferences are the correct explanation for the observed patterns. Accordingly, in this paper we propose an alternative approach that maintains the standard preferences and rational expectations assumptions, yet is focused on how customers react to information they obtain or experiences they go through. Our approach enables us to explore whether simple decision rules can best describe customers' behavior and then to also examine which learning theory is most consistent with these rules. Our findings suggest that consumers rely on *non-standard decision rules* to adjust their behavior, and that, as a result, they may ultimately choose the "wrong", non-cost-minimizing option and pay more after learning. In addition to adopting a different approach to the analysis of individual choice patterns, the current paper also contains analyses that were not included in Ater and Landsman (2013): it examines the *speed and the direction* of learning; it studies the impact of two determinants of switching contracts; it investigates the role of *second-time switching decision*s and customers' learning within a given contract.[1]

The strong association we find between specific triggers and switching behavior also suggests that consumers exhibit *limited attention* or *limited recall*. We are aware of only two empirical studies that focus on the implications of limited attention on individuals' learning (Hanna, Mullainathan and Schwartzstein 2014, Ho, Hogan, and Morton 2015). In both papers, individuals fail to notice, and therefore are unable to learn from, important features of the data they possess. Herein, in contrast, we document instances in which individuals "overreact" to information they possess or to experiences they have undergone (see related models by Bordalo, Gennaioli, and Shliefer 2013. 2015).

In addition to providing empirical evidence to support or refute specific learning theories, our findings contribute to the limited number of studies that use micro-level panel data to examine the consequences of learning and usage-based pricing, typically in the context of a menu of contracts. A common feature of these studies is the assessment of a customer's chosen contract against alternative contracts that are being offered (DellaVigna and Malmendier 2006; Ketcham et al. 2012; Ito 2014; Genakos et al. 2015 Gopalakrishnan et al. 2014; Nevo et al. 2015). A main conclusion in many of these studies is that as consumers gain experience they are able to reduce costs, typically through contract switching. Our work is also related to studies that examined individuals' responses to unexpected events under a given contract.

[1] Unlike Ater and Landsman (2013), this paper only focuses on bank customers who adopted one of the new contracts and exclude from the analysis customers who kept the old pay-per-use pricing scheme.

These studies have shown that, after experiencing an undesirable event, consumers are likely to behave in a way that reduces the likelihood that the event will recur. For instance, Haselhuhn et al. (2012) analyze data from the video rental company 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. Several studies have also examined the impact of personal experiences on financial decisions (Kaustia and Knupfer 2008; Chiang et al. 2011; Choi et al. 2009). Malmendier and Nagel (2011; 2015), for instance, have shown that personal experiences affect financial risk-taking and inflation expectations, respectively. They also find evidence for the recency effect, i.e. that the relevant trigger has a short-term impact on households' financial choices or on the formation of inflation expectations. Finally, other papers on learning use structural models to explore learning patterns (e.g. Lambrecht, Seim and Skiera 2007; Narayanan et al. 2007; Iyengar et al. 2007; Camacho et al. 2011; Goettler and Clay 2011; Grubb and Osborne 2015). In these studies, individuals are typically uncertain about product attributes and/or usage and use Bayesian updating to process signals (e.g., experience or advertising) on product quality and future usage.[2]

The rest of the paper is organized as follows. In Section 2 we describe the main features of the learning theories we consider, and develop testable hypotheses. In Section 3 we describe the data and the market that we study. Section 4 contains the empirical analysis of adoption, contract switching and quitting decisions. In Section 5, we discuss and conclude.

[2] Ching et al. (2013) review empirical structural choice models with learning where individuals are uncertain about product attributes. For a critic on this approach, see Lin et al. (2014).

2. Theoretical Background and Testable Hypotheses

In this section, we review basic features of Bayesian updating, reinforcement learning, and directional learning.[3] We 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.

2.1 Learning Theories

2.1.1 Bayesian updating

Bayesian updating is the prevalent approach used in economics to model individual learning. According to standard Bayesian learning models, over time customers receive information signals that enable them to update their initial beliefs and to learn about the desirability of their choices. As a normative approach, Bayesian updating—in combination with the rational expectations assumption—implies that individuals improve or, at least, do not impair their well-being as they gain more information. This suggests 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.[4]

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 Shliefer 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. In Section 4.4, we return to these modifications of the Bayesian model and argue that they cannot explain our findings. Outside the Bayesian inference framework, researchers have developed alternative learning

[3] Bordalo, Gennaioli, and Shliefer (2013, 2015) propose static theoretical models that emphasize the importance of saliency and customer inattention in consumers' decision making, While their papers do not explicitly model learning, they could also explain our main findings.

[4] Goettler and Clay (2011) present a unique situation where rational expectations might still lead to non-optimal contract choice.

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.

2.1.2 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 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, as compared to a reference payoff, are associated with higher future choice probabilities, whereas lower payoffs, as compared to a reference payoff, 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.

2.1.3 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 also is likely to lead to a correction in a direction that reduces the likelihood of re-experiencing the negative experience.

Despite the basic 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) and, in both theories, individuals assess the 'negativity' of their current experience against a reference point (Erev and Roth 1998, Selten et al. 2005).

2.2 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 on contract switching decisions. We posit that the customers who switch contracts have gone through an active (though potentially distorted) decision-making process, which we consider an indication of learning. Second, in order to investigate the applicability of reinforcement and directional learning theories, we need to identify and characterize the reference payments that customers might use to compare against their actual monthly bills (and subsequently adjust their behavior). On the basis of these reference payments, we can explore the triggers for contract changes, and their effects on customers' eventual actions and payoffs. We consider two different reference payments, which, we argue, customers use to compare against their actual monthly bills. The first reference payment is the monthly fixed payment stipulated in the customer's chosen contract. Specifically, we suggest that a customer compares his actual monthly payment against the fixed payment, and when the actual payment exceeds the latter amount—because of overage fees resulting from usage beyond the contract allowance—he has a negative experience that induces a reevaluation of his past choices. The monthly fixed payment is a natural reference payment in the context of a three-part tariff contract, and previous studies on three-part tariff plans (e.g. Herweg and Mierendorff 2013; Genakos 2015) have used the fixed payment as a reference point.[5] The second reference point we consider is the commission that the customer was used to paying the bank prior to adopting a three-part tariff contract. This pre-adoption payment is a natural comparison for a customer who has just adopted a new payment plan. A large body of evidence suggests that individuals use past personal experience to determine reference payments, and that these references carry disproportional weight in current decisions. For instance, in the housing market, Genovese and Mayer (2001) show how historical sell prices affect future transactions. Accordingly, we consider the ratio between a customer's post-contract adoption payment (including overage payments) and her pre-adoption average monthly payment (the post-pre commission ratio) as a measure for the deviation from the reference payment that the customer expected to pay. A post-pre commission ratio with a value greater than one implies that a customer's mean monthly payment after adoption is higher than her average monthly payment before adoption, and the customer is likely to perceive this experience negatively.

Finally, in order to test the predictions of directional learning, we need to identify an order among the available choices, according to which customers directionally adjust their behavior. Our setting lends itself naturally to this task, as customers are presented with a predefined menu of contracts, ordered according to the size of the allowance, which corresponds to the size of the fixed payments (See Table 1). 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 our setting and the review of the learning theories, we derive five testable hypotheses. The first two hypotheses focus on reinforcement and directional learning theories. According to both theories, behavioral change is driven by negative experiences that customers undergo, and these negative experiences are typically relative to a reference point.

[5] Unlike these studies, we postulate that deviations from the reference payment trigger the customer response but are not part of an individual's utility function.

We further recall that, 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. Based on this discussion, we posit:

Hypothesis 1. According to reinforcement learning theory, customers who deviate from their expected payment and therefore 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.

Within the directional learning theory framework, in contrast to the reinforcement learning framework, it is possible to make a more concrete prediction regarding the choice that a customer will make after a negative experience. In particular, according to directional learning, the contract change will be in a direction that reduces the likelihood of recurrence of the negative experience. For overage payers, an upward switch will reduce the likelihood of paying overage in future periods, and for customers whose payment increases following adoption, a downward switch reduces the likelihood of paying more than the average payment prior to adoption. We thus posit:

Hypothesis 2a. 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 2b. According to directional learning theory, customers who experience a high post-pre commission ratio will switch downward or quit the new subscription contracts altogether.

We now turn to examining the speed of switching as a means of identifying which learning theories might drive customer behavior in our context. Specifically, we investigate whether a customer's switching decisions are influenced by the recency effect, which implies that additional information that individuals gain has a short-lived impact on their subsequent choices. In the case of Bayesian updating, individuals are not predicted to overweigh new information but rather to also rely on prior, or *base-rate*, information when changing their existing choices (e.g., Tversky and Kahneman 1973; Grether 1980). Correspondingly, Bodoh-Creed Benjamin and Rabin (2013) show that if individuals neglect the "Bayesian" base-rate information, a recency effect is observed. Unlike Bayesian updating, both reinforcement and directional learning theories incorporate the recency effect. Thus, we predict:

Hypothesis 3. According to reinforcement and directional learning theories, and in contrast to the expectations of the standard Bayesian approach, individuals are subject to the recency effect (i.e., individuals underweigh prior beliefs and are more likely to react to recent experiences as compared to more distant experiences).

The previous hypotheses focused on the determinants, speed, and direction of contract switching. We now turn to consider the consequences of these decisions. According to the standard Bayesian inference framework with rational expectations, a customer who learns new information is predicted to act in a manner that does not worsen his economic situation. Furthermore, as shown by Rabin (2000) the amount of money that customers lose by switching to a contract with a greater allowance cannot be rationalized by standard risk aversion arguments. In our setting this leads to a natural prediction that customers will only switch to contracts that entail lower payments:

Hypothesis 4. According to Bayesian learning theory with rational expectations, controlling for account usage, customers who switch contracts will not, on average, end up paying more after the switch than they paid before the switch.

Though the standard Bayesian framework cannot explain behaviors that are consistent with the switching to contracts that result in higher payments, modifications to the standard framework may accommodate such phenomena. Specifically, according to Rabin (2002), customers may interpret overage payments as a signal for increased usage and may thus overestimate 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. This is because 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 quit. We thus hypothesize:

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, according to these adjusted Bayesian learning theories, customers who incur overage payments and *expects increased future usage are not expected to return to the pay-per-use pricing scheme, which inherently entails higher payments and greater variation than the three-part tariff contracts in case of increased usage.*

3. Context and Data

3.1 The Menu of Contracts

We use data on the introduction of a menu of six three-part tariff contracts by a large commercial bank that operates in a developed OECD member country. 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 introduction of the new contracts. Under the old pricing scheme, customers paid a commission for each activity they engaged in (e.g. money transfer, cash withdrawal, new deposit). Moreover, in the old pricing scheme, each activity entails a specific price regardless of the channel used (i.e. direct or clerk-assisted). These prices varied substantially across activities, ranging from a few cents to as much as $7 for different activities. 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.[6] 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

[6] Three-part tariff contracts for cellular service also typically include three types of allowances: voice, text and data.

characteristics of the six three-part tariff contracts. Throughout the analysis, the number of the contract is an indication of the size of the allowance and the corresponding fixed payment (e.g., contract 2 has a larger allowance than contract 1 and entails a higher fixed payment). We use this ordering of contracts to analyze the direction of the contract switching decision: upward switching to a contract with a larger allowance versus downward switching to a contract with a smaller allowance.

3.2 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 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 (including pre-adoption months), the data set contained the following information: (i) the contract used for the account during that month; (ii) the number of transactions of each channel (check deposits, direct-channel transactions and clerk-assisted transactions) carried out by the account holder in that month;[7] (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

[7] For each customer and month, we also know how much the customer would have paid under the old-pricing scheme.

levels of savings and loans), and demographic characteristics (e.g., customer age and socio-demographic index[8]). Our data also include the number of direct marketing calls made to each customer prior to adoption 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 the entire research data set.[9] 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 quit the new contracts and return to the old pay-per-use scheme.

3.3 The Market

The bank at the focus of our study 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.

4. 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. We also show that customers who do not switch contracts do not adjust their usage over time in order to fully utilize the quantity of transactions offered under the allowance. Next, we turn to the main analysis and examine how experience with the new contracts affects customers' likelihood of switching as well as the timing of contract switching decisions, and the consequences of these decisions. In the final part of this section we also examine the determinants, timing and consequences of quitting decisions.

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

[9] 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.

4.1 Contract Adoption

In our data, 32,650 customers adopted one of the available three-part tariff contracts. We use information on the set of available contracts and on customers' usage before and after adoption to compare customers' actual contract choices with their 'optimal' contract choices, i.e., contract 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, given his actual usage. Figures 1A and 1B present distributions of the 'distance' associated with customers' chosen contracts, where distance is measured as the number of contracts separating the chosen contract from the optimal contract. For example, if a customer's optimal contract is 2 and he chooses contract 5, then the distance is denoted as '3 contracts above' the optimal contract. The optimal contract for each customer is identified on the basis of his account usage either 3 months before adoption (i.e., ex-ante approach, Figure 1A) or 3 months after adoption (ex-post approach, Figure 1B).[10] We see that the difference between ex-ante and ex-post distributions is very small. This small difference indicates that ex-post consumption behavior cannot explain ex-ante non-optimal contract choices. Specifically, in both figures we observe that only 29% of the adopting customers chose their cost-minimizing contracts. The vast majority of non-optimal contract choices were for contracts with larger allowances than the allowance offered by the cost-minimizing contract. In fact, only about 2% of adopting customers adopted contracts that were 'below' their cost-minimizing contracts (i.e., with lower fixed payments and transaction allowance). In contrast, 25% of adopters chose contracts that were one above their cost-minimizing contracts, 25% chose contracts that were two contracts above their cost-minimizing contracts, and 19% chose contracts that were three or more contracts above their cost-minimizing contracts. Had all customers chosen their cost-minimizing contracts, the average monthly payment across all customers would have been nearly 30% lower.

[10] Changing the relevant time-period before or after adoption has little effect on the fraction of customers who chose their cost-minimizing contracts. In Appendix A we provide more details on the calculations used for the optimality assessment. Note that in the optimality analysis that we present, we do not consider the old pricing scheme as an alternative. That is, the old pricing scheme could not be the optimal option for customers. Once we also include the old pricing scheme as a possible optimal contract, the percentage of customers who chose optimally drops to around 18%. The drop in the fraction of customers who choose the optimal contract results from the fact that for many customers, keeping the old pricing scheme was more cost-effective than the contract they have chosen.

4.2 Post-Adoption Usage Patterns

According to standard learning models, a learning customer who has not chosen her cost-minimizing contract is expected to either switch to a better contract or adjust her consumption behavior in a way that 'justifies' the chosen contract. In particular, a customer who chooses a contract with an excessively large allowance, and does not switch to a lower contract, is expected, over time, to more fully utilize the allowance of 'free' transactions offered by the chosen contract. To examine this conjecture, we focus on customers who retained their initial contract selections throughout our investigated time window, and we compare their usage levels (the number of clerk-assisted and direct transactions, and check deposits) in the first month after adoption with those in the last month of our data. In contrast to what standard learning models would predict, we do not observe an increase in the number of transactions over time among the (many) customers who chose contracts with excessive allowances. Rather, we find that the average usage level in each channel is quite stable over time, and even decreases slightly. The average number of clerk-assisted transactions declined from 0.79 to 0.74 (t-statistic = 5.52), the average number of direct transactions declined from 3.5 to 3.29 (t-statistic = 9.39), and the average number of checks deposited dropped from 3.63 checks to 3.38 (t-statistic = 9.71).

To further examine whether customers who do not switch contracts adjust their consumption behavior to match their chosen contracts, we investigated whether, over time, usage levels became bunched at the various contract allowance levels. According to standard learning models, we would expect to observe a disproportionate concentration of transaction usage levels around the contract allowances over time. Figures 2A and 2B present the distributions of the number of unused transactions in customers' allowances (for check deposits and for clerk-assisted transactions, respectively), in the month after adoption and in the last month of our data.[11] We see that the distributions are similar at the two time points. Specifically, there is no evidence for a higher concentration around zero in the last month of our sample. In other words, we do not observe bunching of usage levels around the transaction

[11] Note that we do not investigate direct transactions here since in all but one contract the allowance for direct transaction is unlimited. Similarly, the analysis of clerk-assisted activities in Figure 2B focuses on contracts 3-5.

allowances as customers gain experience with the contracts. This absence of bunching, together with the fact that customers do not switch contracts despite these unused allowance transactions (a phenomenon also documented in other studies, e.g., Grubb and Osborne 2015), suggests that the customers we analyze are inattentive to the remaining balance of their contracts and do not correct non-optimal contract choices by changing their consumption behavior. As we document in the subsequent section, consumers seem to be more attentive to the gap between the expected (i.e., reference) payment level and their actual monthly payments.

4.3 Contract Switching

In this section we focus on customers who adopted three-part tariff contracts and later switched to a different three-part tariff contract. 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 important given customers' strong tendency to initially adopt contracts with larger than optimal allowances.

4.3.1 Determinants and speed of switching

4.2.1.1 Descriptive statistics

Table 2 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 2, 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 the post-pre commission ratio is the ratio between payments after contract adoption (and before switching/quitting if they take place) and payments before contract adoption. Thus, higher values are likely to reflect a more negative experience for the customer. Figure 3A plots the distribution of this variable for all adopting customers (mean of 1.17, std. 0.45). Figure 3B plots the distribution of the mean post-pre commission ratio for switching customers, grouped according to the direction of their switches. We see that customers who switched downward exhibit a higher post-pre commission ratio, on average, compared with customers who switched upward (mean of 1.30 vs. 1.12, for downward- and upward-switchers, respectively).

Overall, the descriptive findings seem consistent with the idea that the fixed contract payment and the pre-adoption monthly payment indeed serve as reference payments, and that when customers undergo the negative experience of exceeding one of these reference payments, they tend to switch to contracts in which the negative events they have experienced are less likely to recur (Hypothesis 2a and 2b).

To obtain an initial idea of whether the influence of these negative experiences is subject to a recency effect, we compare the elapsed time until the switch for downward- and upward-switchers. We expect that contract switches driven by high post-pre commission ratios will occur sooner after adoption compared with contract switches triggered by overage payments. This is because customers are likely to realize that they made an ill-suited adoption choice shortly after adoption, whereas a customer experiences overage payments only after exceeding his or her allowance, and not necessarily right after adoption. Figure 4 plots the cumulative distribution of elapsed time to switch, divided into downward- and upward-switchers. 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 5 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. Thus, downward contract switches occur shortly after adoption, i.e., shortly after the experience of a high post-pre commission ratio, and upward contract switches occur shortly after overage payments. Figure 6 provides similar evidence with regard to quitting decisions. The Figure displays the elapsed time between adoption and quitting decisions for customers who paid overage payments and those who have not paid overage payments before quitting. Like with switching customers, the figure demonstrates that quitters who have not paid overage payments quit before those who did pay overage payments. Arguably, this is because many customers experience overage payments soon after they realize their post-pre commission ratio. To further support these descriptive findings, we now turn to regression analyses.

4.3.1.2 Identifying the triggers for switching

We are interested in investigating the role of overage payments and post-pre commission ratio as triggers for contract switching (Hypotheses 1 and 2). 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:

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

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:

$$(2)\, Z_i(t) = \{PostPre\ Ratio_{\ it-1}, Overage_{it-1}, \overline{Overage_{it-2}}, Transact_{it-1}, X_{it-1}\}$$

For each post-adoption period prior to switching, the main explanatory variables are customer *i's* post-pre commission ratio at time t-1, $PostPreRatio_{i,t\text{-}1}$, and the amount paid in overage fees at time t-1, $Overage_{i,t\text{-}1}$. Post-pre commission ratio 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 payment. In addition, $\overline{Overage_{i,t-2}}$ represents the mean monthly amount paid prior to t-1.[12] We also control for the number of transactions of the customer. $Transact_{it}$ is a matrix that includes the number of clerk-assisted and direct transactions and the number of checks that customer i deposited during 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 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.[13]

The results, presented in Table 3, 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 post-pre commission ratio we find the opposite. Post-pre commission ratio 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.[14] Thus, overall, our findings are consistent with both the prediction of reinforcement learning theory regarding the higher likelihood to switch after experiencing a negative experience (Hypothesis 1), and the predictions of directional learning theory, which also include the direction of the switching decision (Hypotheses 2a and 2b). In order to examine Hypothesis 3, regarding the longevity of trigger influence over time, we compared the coefficient for the immediate influence of overage payments with the coefficient for their influence over longer periods of time. That is, we included in the hazard regression both the amount paid as overage

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

[13] We analyze quitting decisions in section 4.3.

[14] Consistent with these results, additional panel data regression results (available upon request) show that downward-switchers paid significantly higher payments after they adopted a three-part tariff contract compared to their pre-adoption commission payments. In contrast, customers who adopted a three-part tariff contract and later switched upward did not pay higher payments after they adopted a three-part tariff contract (yet before their subsequent switch). These results further suggest that non-optimal contract adoption decisions are associated with subsequent contract switching.

at time t-1, and the mean monthly amount paid prior to t-1 (i.e., from adoption up to 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 reinforcement and directional learning (Hypothesis 3).

4.3.2 Consequences of upward- and downward-switching

We now turn to examine the consequences of switching (Hypothesis 4). Hypothesis 4 enables us to explore the joint assumption of rational expectations and Bayesian learning. To test this hypothesis, we exploit the longitudinal nature of our data and estimate the following panel data fixed effects regression:

$$(3)\ Log(payment)_{it} = \beta_0 + \beta_{0i} + \beta_{0t} + \beta_1 PostSwitch_{it} + \beta_T Log(Transact_{it} + 1) + \beta_X Log(X_{it} + 1) + \varepsilon_{it}$$

where the variable $PostSwitch_{it}$ in Equation (3) is a dummy variable that equals one if customer i has switched contracts by time t and zero otherwise. $Transact_{it}$ and X_{it} are as defined above. We implement log-transformation for all the variables that are not binary variables. Finally, we also include account (β_{0i}) and time (β_{0t}) fixed effects to control for unobserved differences across customers, and unobserved time trends. The standard errors at the individual account level are clustered. The regression results are reported in Table 4. In columns 1 and 2 we present the regression results for 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 upward increased by 11 percent after switching, thus contradicting Hypothesis 4. We find additional support for these patterns when we restrict attention to customers who switched contracts twice. We estimate Equation (3) for customers who switched contracts twice, considering only the months after the first switch.[15] The results, reported in columns 3 and 4 in Table 4, are qualitatively similar to the results presented in columns 1 and 2, for the first switch. Customers whose second switch was a downward switch reduced their monthly payments by 27.5 percent, while customers whose second switch was an upward switch increased their monthly payments by 7.4 percent.

These findings support our proposition that downward switchers, on the one hand, aim to avoid re-experiencing a high post-pre commission ratio. 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).[16] Upward switchers, on the other hand, do not seem to focus on the post-pre commission ratio. These customers' switching decisions are triggered by overage payments, and thus their new contracts are contracts that limit the potential for this negative experience to recur. 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 (averages of $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 recency effect, are consistent with directional learning models, yet are difficult to explain within the standard Bayesian learning framework.

4.4 Quitting Decisions

We now turn to the analysis of quitting decisions and their implications. This analysis is necessary to test Hypothesis 5 and rule out modifications to the standard Bayesian updating framework. Specifically, we investigate how overage payments and experiences of high post-pre commission ratio affect the tendency of customers to quit the new subscription contracts in

[15] 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.
[16] These estimates are based on a regression analysis that is available upon request from the authors (see footnote 13).

favor of the old pay-per-use scheme. We investigate how each trigger affects quitting by estimating the following hazard regression:

$$(4)\, h_{q,it} = h_{q0}(t) \cdot e^{\beta_q' Z_i(t)}, \qquad i = 1 \dots n$$

The dependent variable in Equation (4) 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 (2).

Table 5 presents the regression results of the hazard regression for quitting. Higher overage payments in the month prior to quitting and lower post-pre commission ratio are 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 3), 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).

4.5 Summary

Overall, our empirical analysis provides evidence that supports Hypotheses 1, 2, and 3, yet contradicts Hypotheses 4 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 3). We also find that high post-pre commission ratio and overage payments trigger contract switches (Hypotheses 1 and 2). The type of trigger experienced seems to influence the direction of the switch: compared with customers who experienced overage payments, customers with high post-pre commission ratio were more likely to switch to contracts with lower fixed payments, and were much less likely to switch to contracts with higher fixed payments (Hypothesis 2). We also find that the majority of customers, and specifically those who switched upward, ended up paying more to the bank after switching, thus contradicting Hypothesis 4. Conversely, we see indications for contract switching to contracts that eliminate the probability of re-experiencing the triggers to switching. Finally, we find that overage payments also trigger quitting decisions, a phenomenon that cannot be explained by customers' overestimation of their future consumption (Hypothesis 5).

5. 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 the initial adoption decision, post-adoption consumption behavior, and subsequent contract changes (or lack thereof). 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, according to the standard Bayesian approach, learning in our setting can take two forms. First, a customer may learn from post-adoption experience and change her consumption behavior in a way that makes her chosen contract optimal. However, altering one's consumption behavior may not always be a desirable or feasible option. In these cases, we expect the customer to utilize the information he gains after adoption and switch to the contract that is optimal for him. Our data suggest that neither form of learning takes place as expected—most adopting customers neither switch to new contracts nor change their usage, even when their chosen contracts are non-optimal. Those who do switch to new contracts are triggered by specific experiences and, contrary to what Bayesian learning models would predict, select non-optimal contracts; that is, they may end up paying *higher* monthly commissions than they paid before switching. To explain these findings we propose that consumers use simple decision rules and that these decision rules are most consistent with directional learning theory, and are generally inconsistent with Bayesian learning models. Notably, we are the first to provide empirical evidence using real-life choices that support the directional learning theory. Our findings are particularly striking given the simple structure of the environment we study. Previous empirical settings typically involved several potential confounding characteristics (e.g., new service with high levels of customer uncertainty). Checking accounts, on the other hand, 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 other customers. Accordingly, the failure of Bayesian updating to explain our findings does not hinge on the difficulty to form correct beliefs over other players' strategies.

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, suggesting 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., Driscoll, J. C., Gabaix, X., & Laibson, D. I. (2013). Learning in the credit card market. Mimeo, Harvard University.

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

Bodoh-Creed, A., Benjamin, D., & Rabin, M. (2013). *The Dynamics of Base-rate Neglect*. Mimeo Haas Business School.

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

Bordalo, P., Gennaioli, N., & Shleifer, A. (2015). Memory, attention and choice, unpublished working paper.

Busse, M. R., Pope, D. G., Pope, J. C., Silva-Risso, J. (2015). The Psychological Effect of Weather on Car Purchases. The Quarterly Journal of Economics, 130(1), 371-414.

Camacho N, B. Donkers, S. Stremersch (2011). Predictably non-Bayesian: Quantifying salience effects in physician learning about drug quality. *Marketing Science*, 30(2):305–320.

Chiang, Y. M., Hirshleifer, D., Qian, Y., & Sherman, A. E. (2011). Do investors learn from experience? Evidence from frequent IPO investors. *Review of Financial Studies*, hhq151.

Ching, A, Erdem T, & Keane M. 2013. Learning Models: An Assessment of Progress, Challenges and New Developments. *Marketing Science,* 32(6), 913–938.

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., 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.

Genakos, C., Roumanias, C., & Valletti, T. 2015. Loss aversion on the phone. Working paper, ECARES.

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

Genovese, David, and Christopher Mayer. (2001). Loss Aversion and Seller Behavior: Evidence from the Housing Market. *Quarterly Journal of Economics*, 116(4): 1233–60

Goettler, R. L., & Clay, K. (2011). Tariff choice with consumer learning and switching costs. *Journal of Marketing Research*, *48*(4), 633-652.

Gopalakrishnan, A., R. Iyengar, and RJ. Meyer. (2014). Consumer Dynamic Usage Allocation and Learning Under Multipart Tariffs. Marketing Science, 34(1), 116–133.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics*, 95(3), 537-557.

Grubb, M. D. (2015a). Overconfident consumers in the marketplace. *Journal of Economic Perspectives*, 29(4), 9-36.

Grubb, M. D. (2015b). Behavioral Consumers in Industrial Organization. *Review of Industrial Organization* (forthcoming).
Grubb, M.D., M. Osborne. 2015. Cellular Service Demand: Biased Beliefs, Learning and Bill-shock. *American Economic Review, 105:1*, 234-271.
Hanna, R., Mullainathan, S., & Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *The Quarterly Journal of Economics*, *129*(3), 1311-1353.
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.
Ho, K., Hogan, J., & Morton, F. S. (2015). *The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program* (No. w21028). National Bureau of Economic Research.
Hogarth, R.M., H.J. Einhorn. (1992). Order Effects in Belief Updating: The belief-adjustment Model. Cognitive Psychology, 24, 1–55.
Iyengar R, Ansari A, Gupta S. (2007). A Model of Consumer Learning for Service Quality and Usage. *Journal of Marketing Research*, 44(4), 529–544.
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, M., & Knüpfer, S. (2008). Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions. *The Journal of Finance*, *63*(6), 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.
Lin, S., Zhang, J., & Hauser, J. R. (2014). Learning from experience, simply. *Marketing Science*, *34*(1), 1-19.
Malmendier, U., S. Nagel. (2011), Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking? Quarterly Journal of Economics, 126, 373-414
Malmendier, U., & Nagel, S. (2015). Learning from inflation experiences. *Unpublished manuscript, UC Berkley*.
Narayanan S, Chintagunta P, Miravete E. (2007). The Role of Self Selection, Usage Uncertainty, and Learning in the Demand for Local Telephone Service. *Quantitative Marketing and Economics*, 5(1), 1–34.
Nevo, A., Turner, J. L., & Williams, J. W. (2015). Usage-based pricing and demand for residential broadband. *Available at SSRN 2330426*.

Ockenfels, A., R. Selten. (2014). Impulse balance in the newsvendor game, *Games and Economic Behavior*, 86, July, 237–247.
Ortoleva, P. (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, M. **(**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., & Stoecker, R. (1986). End behavior in sequences of finite Prisoner's Dilemma supergames A learning theory approach. *Journal of Economic Behavior & Organization*, *7*(1), 47-70.
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
Simon, H. A. (1959). Theories of decision-making in economics and behavioral science.
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., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological bulletin*, *76*(2), 105.
Tversky, A., D. Kahneman. (1973). Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology*, 5(2), 207-32.
Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. J. Risk Uncertainty 5(4):297–323.
Woodford, M. (2003). Interest and Prices: Foundations of a Theory of Monetary Policy.: Princeton University Press.

TABLE 1 –MENU OF THREE-PART TARIFF CONTRACTS17

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
2	$6.25	$1.2	$0.30	0	Unlimited	10
3	$6.75	$1.2	$0.30	3	Unlimited	7
4	$7.75	$1.2	$0.30	3	Unlimited	10
5	$9.50	$1.2	$0.30	7	Unlimited	12
6	$14	$1.2	$0.30	Unlimited	Unlimited	15

[17] In the month when the new contracts were introduced, bank customers could choose from a menu of four three-part tariff contracts (contracts 1, 3, 5 and 6 in Table 1). Nine months after the introduction of these four contracts, contracts 2 and 4 were added to the set of available contracts while contract 3 was removed. In addition, the allowance of direct transactions in contract 1 was reduced from unlimited to 7 (however, customers who chose these contracts before they were removed were entitled to keep their terms afterwards). Our optimality analysis takes these modifications into account.

TABLE 2 – 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	2030

[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 3 – HAZARD REGRESSION ANALYSIS FOR SWITCHING DECISION

Variable	Downward-Switching Hazard Regression: Parameter Estimate (Standard Error)	Downward-Switching Hazard Regression: Hazard ratio	Upward-Switching Hazard Regression: Parameter Estimate (Standard Error)	Upward-Switching Hazard Regression: Hazard ratio
Post-pre commission ratio[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	

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

[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

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 4 – CUSTOMER SWITCHING PAYMENT REGRESSIONS

VARIABLES	*Upward switchers (over post-adoption months)*	*Downward switchers (over post-adoption months)*	**'Second-time' upward switchers**	**'Second-time' downward switchers**
Post switching month	0.119***	-0.325***	0.074***	-0.275***
	(0.004)	(0.008)	(0.023)	(0.041)
Number of clerk-assisted transactions	0.152***	0.093***	0.091***	0.114***
	(0.004)	(0.006)	(0.019)	(0.037)
Number of direct transactions	0.010***	0.005	0.024	0.005
	(0.002)	(0.003)	(0.017)	(0.004)
Number of check deposits	0.056***	0.044***	0.092***	0.045***
	(0.004)	(0.009)	(0.026)	(0.010)
Number of owners	0.039	0.135**	-0.024	0.156**
	(0.036)	(0.059)	(0.082)	(0.063)
Parental Social Security benefits	0.006	-0.203*	-0.117	-0.194*
	(0.047)	(0.108)	(0.162)	(0.108)
Elderly Social Security benefits	0.035	0.023	-0.321	0.021
	(0.026)	(0.036)	(0.981)	(0.037)
Number of information inquiries	0.006***	0.010***	0.012	0.009***
	(0.002)	(0.002)	(0.012)	(0.003)
Salary	0.003	-0.003	0.008	-0.008
	(0.007)	(0.013)	(0.022)	(0.013)
Number of salaries	0.015*	0.012	0.036	0.012
	(0.008)	(0.011)	(0.035)	(0.014)
Loans	0.004	0.018***	-0.012	0.022***
	(0.004)	(0.007)	(0.018)	(0.008)
Savings	0.015***	0.003	0.021	0.003
	(0.003)	(0.005)	(0.013)	(0.006)
Constant	1.848***	2.084***	1.791***	2.079***
	(0.020)	(0.037)	(0.089)	(0.037)
Observations	24,379	11,627	987	585
R-squared	0.421	0.364	0.223	0.353
Number of customers	1,441	827	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.
In columns 1 and 2 we focus on upward and downward switchers, respectively, and the estimation results include only the months after contract adoption (Equation (4) in the text).
In columns 3 and 4, the estimation focuses on customers who switched contracts twice and include only the months after the first switching decision. The sample of customers shown in column 3 only includes customers whose second switch was to a contract with larger allowance, while in column 4 we focus on customers whose second switch was to a contract with a smaller allowance.

TABLE 5 – HAZARD REGRESSION ANALYSIS FOR QUITTING DECISION

	Parameter Estimate (Standard Error)	***Hazard ratio***
Post-pre commission ratio[a]	0.02*** (0.00)	1.02
Overage[a]	0.06*** (0.00)	1.06
Mean Past Overage[a]	0.01 (0.01)	1.01
Loans[a]	-0.03*** (0.01)	0.97
Parental Social Security benefits (for children below the age of 18) [a]	-1.63*** (0.40)	0.20
Elderly Social Security benefits[a]	-0.01 (0.11)	0.99
Monthly number of clerk-assisted transactions[a]	-0.12*** (0.02)	0.89
Monthly number of direct transactions[a]	-0.13*** (0.01)	0.88
Monthly number of check transactions[a]	-0.06*** (0.01)	0.95
Number of owners[a]	-0.07 (0.05)	0.93
Number of salaries [a]	-0.18*** (0.04)	0.83
Savings[a]	0.00*** (0.00)	1.00
Salary[a]	-0.11*** (0.04)	0.89
Monthly number of account status inquiries[a]	-5.83E-04 (0.00)	1.00
Account tenure	-2.13E-03 (0.00)	1.00
Socio-economic indicator(scale of 1 to 10)	0.02 (0.01)	1.02
Pre adoption number of marketing calls	0.08*** (0.02)	1.08
Customer age	3.90E-03** (0.00)	1.00
Number of events	2,030	

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

[a]Time-varying variable

[b] Non-time-varying variable

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 – EX-ANTE DISTRIBUTION: DISTANCE OF CHOSEN CONTRACT FROM OPTIMAL CONTRACT

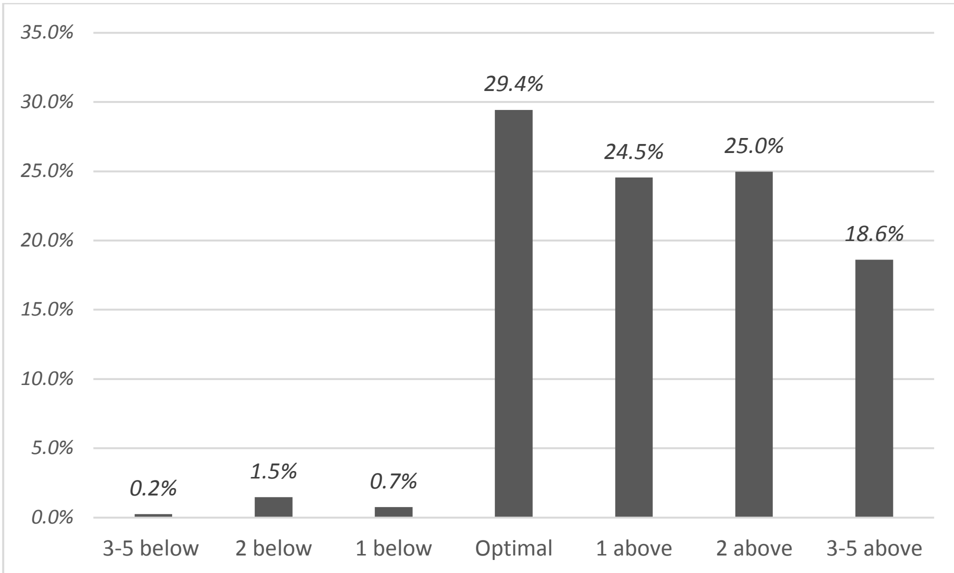

FIGURE 1B – EX-POST DISTRIBUTION: DISTANCE OF CHOSEN CONTRACT FROM OPTIMAL CONTRACT

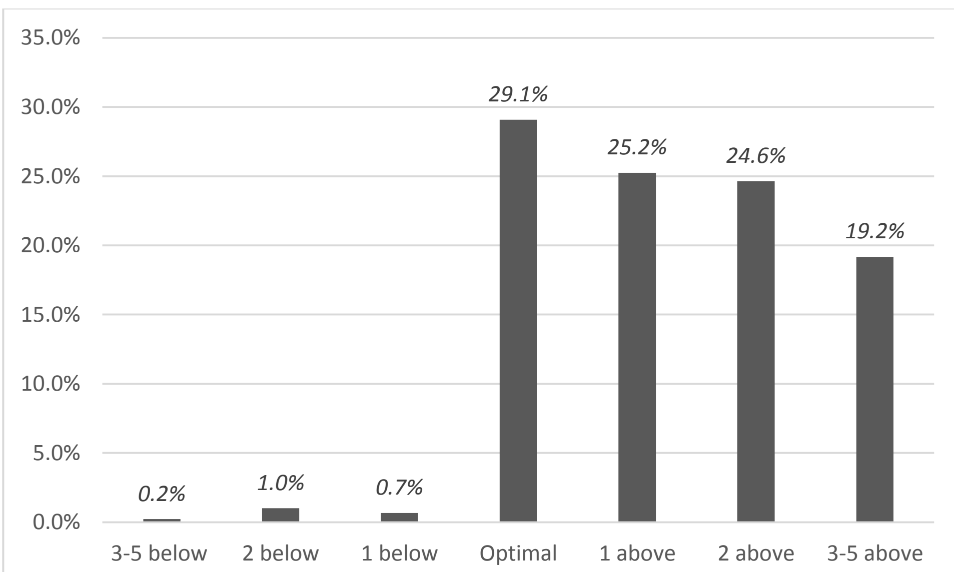

Figures 1A and 1B present the distributions of the 'distance' between chosen and optimal contracts for the individual customer (i.e., the number of contracts separating the chosen contract from the optimal contract for each customer). For example, according to Figure 1A, for 25% of the customers this distance is 2 contracts above the optimal contract. The optimal contract for each customer is identified on the basis of his account usage either 3 months before adoption (1A), or 3 months after adoption (1B).

FIGURE 2A – DISTRIBUTIONS OF UNUSED CHECK DEPOSIT ALLOWANCE

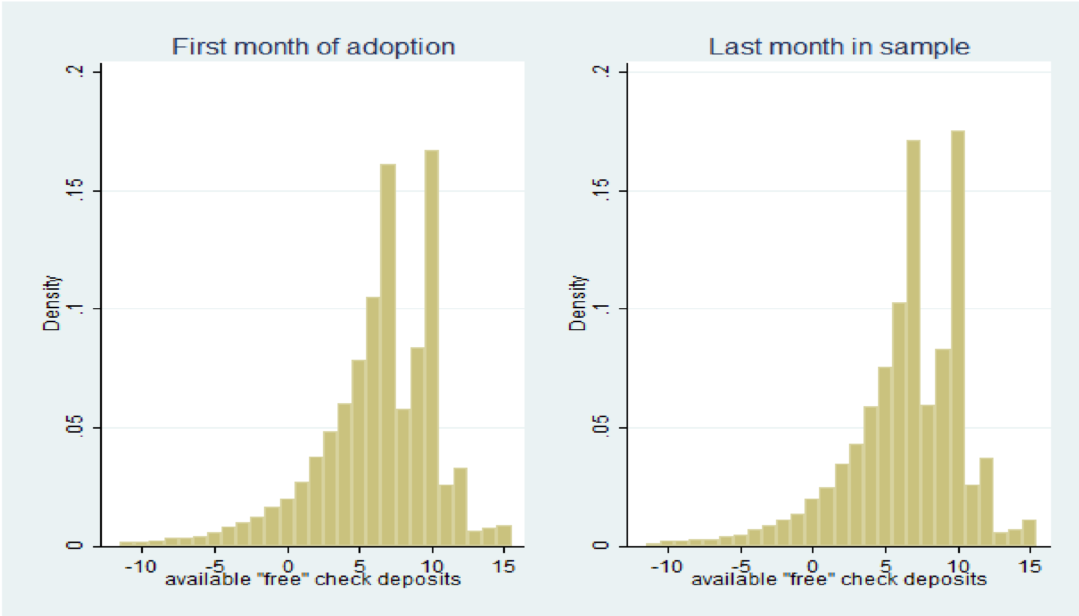

FIGURE 2B – DISTRIBUTIONS OF UNUSED CLERK-ASSISTED TRANSACTION ALLOWANCE

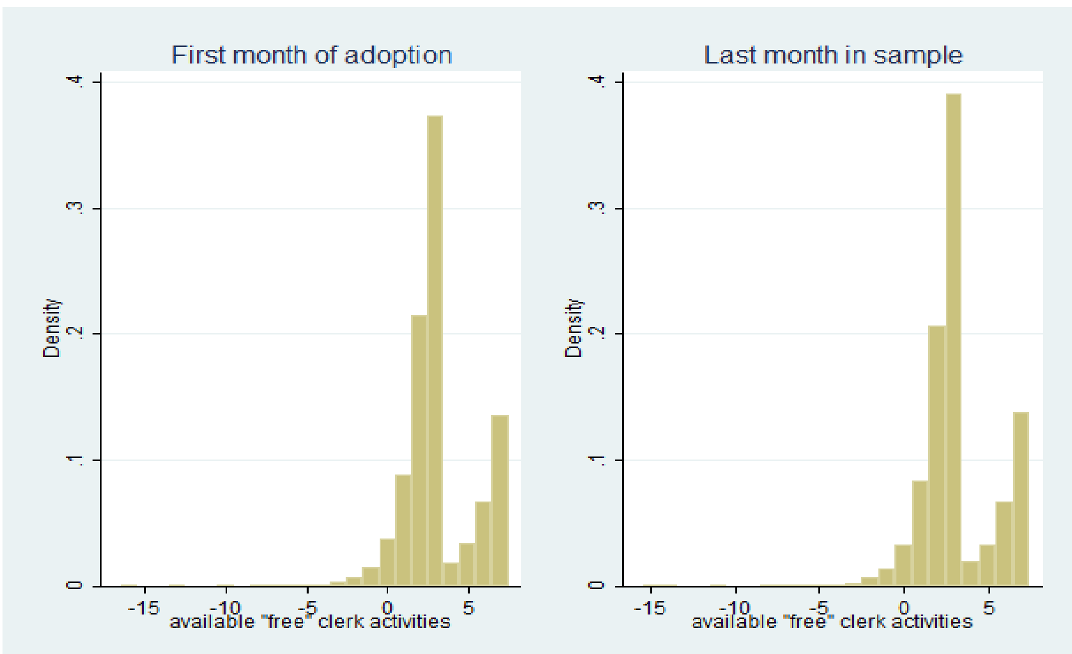

Figures 2A and 2B present the distributions of the number of unused transactions in customers' allowances (for check deposits and for clerk-assisted transactions, respectively), in the month after adoption and in the last month of our data. The horizontal axis is calculated as the actual usage of a customer in the given month minus the contract allowance of the contract of that customer in that month. A positive number reflects a state of unused transactions, while a negative number reflects a state of usage that exceeds the contract allowance. The distribution is for customers who neither switched nor quitted and whose contracts include finite allowances. Both Figures show considerable mass on unused transactions. There is no observable difference in usage between the first and the last month.

FIGURE 3A – POST-PRE COMMISSION RATIO AMONG ALL ADOPTERS

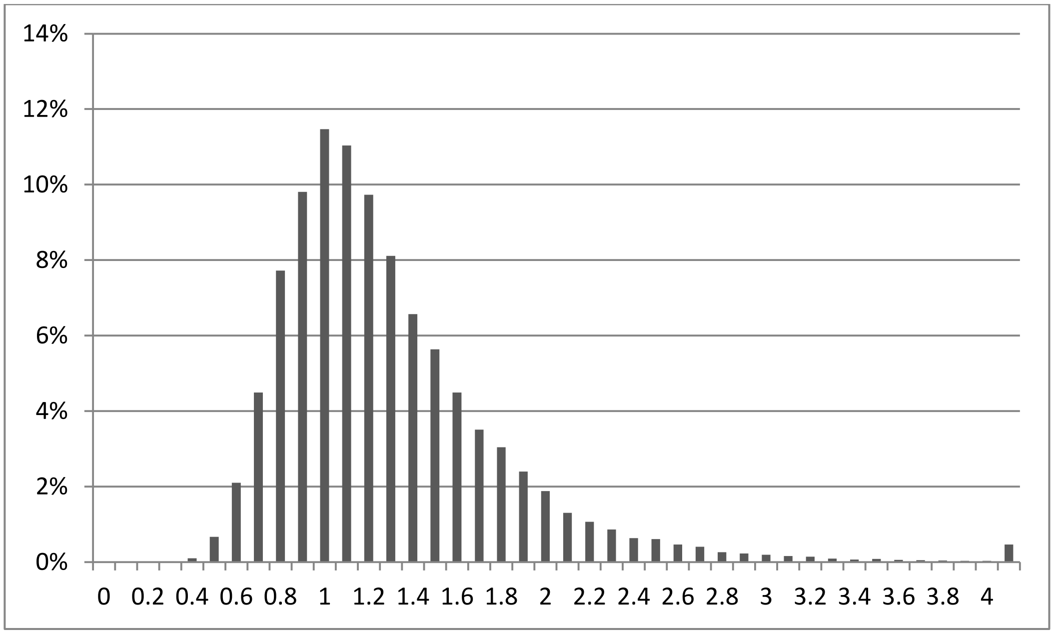

FIGURE 3B – POST-PRE COMMISSION RATIO AMONG UPWARD- AND DOWNWARD-SWITCHERS

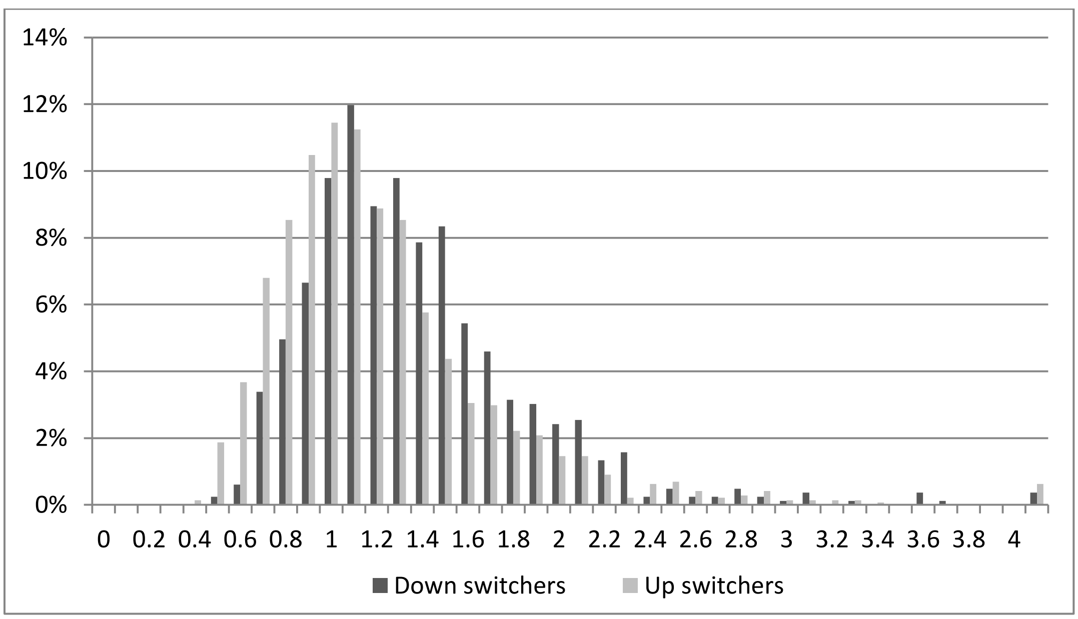

Figures 3A and 3B present the distribution of the mean post-pre commission ratio measure across individual customers (i.e., the ratio between payments after contract adoption and payments before contract adoption). Figure 3A plots this distribution for all adopting customers. Figure 3B plots this distribution for switching customers, grouped according to the direction of their switches. Customers who switched downward exhibit a higher post-pre commission ratio, on average, compared with customers who switched upward, reflecting a more negative experience for downward switchers.

FIGURE 4– CUMULATIVE DISTRIBUTION OF TIME UNTIL SWITCH

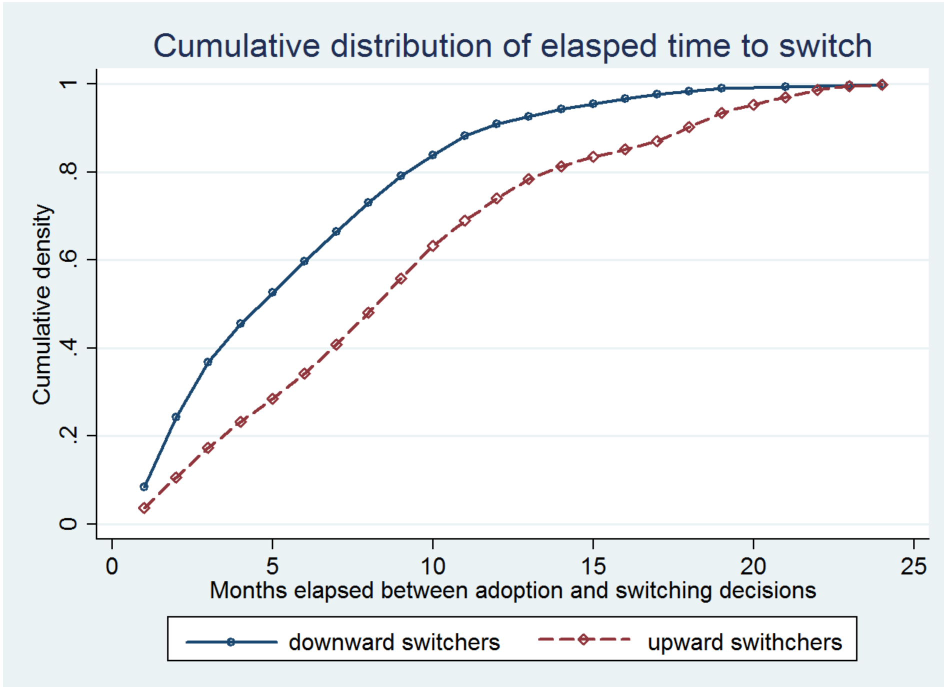

Figure 4 presents the cumulative distribution of elapsed time to switch, divided into downward (blue solid line)- and upward (red dashed line) -switchers. On average, downward switching decisions occur sooner after adoption compared with upward switching decisions.

FIGURE 5 – HISTOGRAM OF TIME BETWEEN LAST OVERAGE PAYMENT AND SWITCH

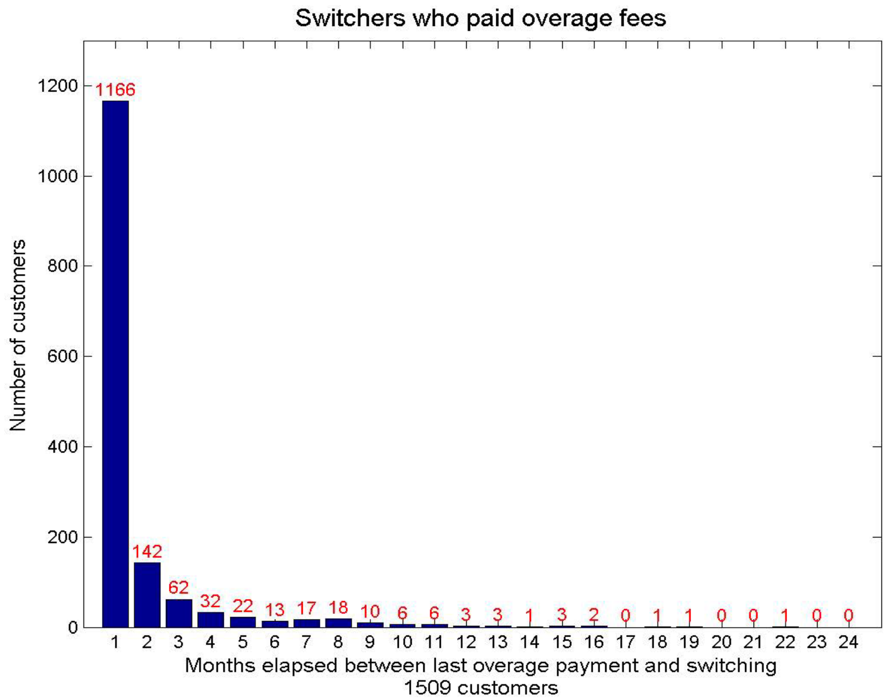

Figure 5 presents the distribution of the elapsed number of months between the last overage payment for a given customer and her subsequent switching decision. The sample of customers includes 1509 customers who paid overage payments before they switched. Of these customers, 1,166 paid overage payments in the month before switching, while only 32 customers did not pay overage payments in the 3 months before switching.

FIGURE 6 – CUMULATIVE DISTRIBUTION OF TIME UNTIL QUITTING

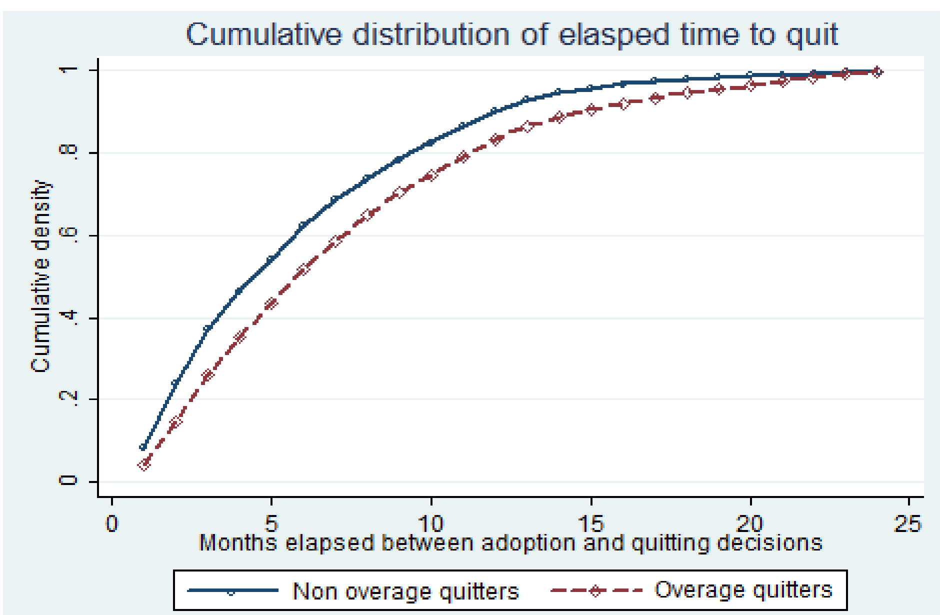

Figure 5 presents the cumulative distribution of elapsed time to quit, divided into overage payers (blue solid line)- and non-overage payers (red dashed line) who quit. On average, the quitting decisions of non-overage payers occur sooner after adoption compared with those for overage payers.

Online Appendix A – Optimality Calculation Schemes

A customer might calculate the cost of a contract in a given three-part tariff contract menu in several ways. For example, customers might consider only their activity in the month of choice while evaluating the contracts. 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 contracts based on their activity in each month, or alternatively calculate the payment associated with each contract 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 contracts, 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 contract. 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 contract based on the customer's actual usage over the relevant timeframe. While the latter approach is more accurate in terms of contract 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 contract 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 contract and paying larger amounts each month. Third, a customer might choose a contract 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' contract 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 contract 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 contract 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 contract payment using the overall payment calculation basis for the four contracts that were available in those three months, and for the old pay-per-use payment system. According to Table A3, Contract 1 is the optimal contract because the overall payment for this contract is the lowest.

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

	Contract 1	Contract 2	Contract 3	Contract 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 contract payment using the mean monthly activity level calculation scheme. According to this calculation scheme the optimal contract is Contract 3.

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

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