

Try It, You'll Like It-Or Will You? The Perils of Early Free-Trial Promotions for High-Tech Service Adoption

Citation for published version (APA):

Foubert, B., & Gijsbrechts, E. (2016). Try It, You'll Like It-Or Will You? The Perils of Early Free-Trial Promotions for High-Tech Service Adoption. *Marketing Science*, 35(5), 810-826.
<https://doi.org/10.1287/mksc.2015.0973>

Document status and date:

Published: 01/01/2016

DOI:

[10.1287/mksc.2015.0973](https://doi.org/10.1287/mksc.2015.0973)

Document Version:

Accepted author manuscript (Peer reviewed / editorial board version)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Try It, You'll Like It—Or Will You?

The Perils of Early Free-Trial Promotions for High-Tech Service Adoption

Bram Foubert

Els Gijsbrechts*

* Bram Foubert (corresponding author) is an assistant professor at Maastricht University, School of Business and Economics, Department of Marketing and Supply Chain Management, P.O. Box 616, 6200 MD Maastricht, The Netherlands, phone: +31 43 3883816, fax: +31 43 3884918, e-mail: b.foubert@maastrichtuniversity.nl. Els Gijsbrechts is Professor of Marketing at Tilburg University, Tilburg School of Economics and Management, Department of Marketing, P.O. Box 90153, 5000 LE Tilburg, The Netherlands, phone: +31 13 4668224, fax: +31 13 4668354, e-mail: e.gijsbrechts@uvt.nl. We thank Charlotte Rolef for her help with the data collection and Kathleen Cleeren and George Knox for their valuable comments. Furthermore, we gratefully acknowledge the many constructive suggestions from the editor, the area editor, and the two reviewers.

Try It, You'll Like It—Or Will You?

The Perils of Early Free-Trial Promotions for High-Tech Service Adoption

Abstract

The proliferation of free trials for high-tech services calls for a careful study of their effectiveness, and the drivers thereof. On the one hand, free trials can generate new paying subscribers, by allowing consumers to become acquainted with the service free of charge. On the other hand, a disappointing trial experience might alienate potential customers, when they decide not to adopt the system and are lost for good. This dilemma is particularly worrisome in early periods, when service quality has not been “tried and tested” in the field, and breakdowns occur. We accommodate these phenomena in a model of consumers’ free-trial and regular adoption decisions. Among other effects, it incorporates usage- and word-of-mouth-based learning about quality in a setting where quality itself is evolving. Consumers are forward-looking in that they account for changes in quality and anticipate uncertainty reduction due to trial usage. We estimate our model and run simulations on the basis of a rich and unique data set that incorporates customers’ trial subscription, adoption, and usage behavior for an interactive digital television service. The results underscore that free trials constitute a double-edged sword, and that timing and consumers’ usage intensity during the trial are key to the effectiveness of these promotions. Implications for managers are discussed.

Keywords: free-trial promotions, adoption behavior, high-tech consumer services, contractual services, learning, promotion effectiveness, promotion timing, usage

1. Introduction

With the rise of high-tech consumer services involving contractual arrangements, free-trial promotions have gained widespread acceptance—becoming the rule rather than the exception. Examples particularly abound for software packages (e.g., antivirus software), communication (e.g., telephone, Internet), and entertainment services (e.g., interactive television). Free-trial promotions offer consumers free access to a full-fledged or restricted version of the service for a limited period. A popular example is the online movie rental service Netflix, which heavily advertises its two- or four-week free-trial offer (Chatterjee, Carroll, & Spencer, 2009). After the trial, consumers can opt out or do nothing and automatically become paying subscribers.

Despite their pervasive presence, free-trial promotions are not well understood. Netflix boasts a trial-to-adoption conversion rate of approximately 90% (Chatterjee et al., 2009), but many post-trial adopters might have adopted anyway, even without the trial. Moreover, the 10% of nonadopters could include people who would have signed a regular contract in the absence of a free trial but decided not to adopt after having tried the system at a moment that service quality was subpar. This concern is particularly relevant for high-tech services, whose quality may change over time due to fluctuations in technical service performance (e.g., Zhang & Niu, 2014). Yet, surprisingly little research considers how the timing of a free trial and consumers' usage intensity during the trial might affect adoption of high-tech consumer service contracts.

Although scarce compared with the vast research on price promotion, studies of free-trial and sample promotions offer some valuable insights. Early research indicates that samples exert lasting effects on sales (Heiman, McWilliams, Shen, & Zilberman, 2001). Bawa and Shoemaker (2004) identify different phenomena underlying a sample's sales impact. A sample may initially cannibalize regular brand purchases from consumers who would have been willing to *buy* the

product. However, after the promotion, the brand likely benefits from accelerated purchases (by consumers who adopt sooner than they would have without the sample) and market expansion (buyers who would not have adopted the product without the sample). Existing research also addresses *why* consumers are more likely to adopt after a sample or trial promotion. Halbheer, Stahl, Koenigsberg, and Lehmann (2014) show analytically that the effectiveness of a (permanent) free trial for information goods depends on the trial's impact on the valence of consumers' quality beliefs;¹ Marks and Kamins (1988) find that samples reduce belief uncertainty. Apart from these direct effects, a trial can influence adoption also indirectly, through word of mouth (WOM). People who have tried a product may initiate brand-related conversations and alter other consumers' beliefs. The more people engage in WOM, the more effective sampling becomes (Jain, Mahajan, & Muller, 1995).

These findings are insightful, but the unique characteristics of our focal setting—contractual high-tech services—suggest some important differences. First, the quality of high-tech consumer services is not stationary but fluctuates over time (e.g., Zhang & Niu, 2014), such that the valence of the trial experience and the trial's WOM effects depend on when the promotion is offered. While quality fluctuations are intrinsic to services in general, they are particularly common in high-tech service settings. In pursuit of early-mover advantages, high-tech service companies may trade off quality to speed up market entry and lock in a substantial number of customers. As a result, quality may be subpar and breakdowns may occur especially early in the service's lifecycle. Extant sampling and free-trial literature suggests that allowing consumers to try a product or service at least does not *reduce* the number of adopters; that is, even if it does not help, it will not hurt (e.g., Bawa & Shoemaker, 2004; Heiman et al., 2001; Jain

¹ In a related context, Pauwels and Weiss (2008) examine the implications when an online content provider moves from “free to fee.” These investigations differ from our setting in not only the focal context but also the studied instruments (permanent pricing instead of temporary promotional offers).

et al., 1995). In contrast, we expect that if the service's quality is insufficient, the trial could lead to a substantial loss of customers who, absent the possibility to first try the service, would have entered into a paid long-term contract.

Second, a free service trial does not involve a fixed consumption *amount* but rather a fixed consumption *period*. Experience accumulation and learning thus is not fixed across triers but depends on their usage intensity. Most sampling studies consider consumer packaged goods, for which sample acceptance and consumption largely coincide (e.g., Bawa & Shoemaker, 2004). The few studies that consider trials of durable products, services, or information goods ignore the effects of usage on trial performance (e.g., Halbheer et al., 2014; Jain et al., 1995; Pauwels & Weiss, 2008). Yet, understanding the impact of usage is essential in these settings: it helps managers assess the likely outcome of an ongoing trial promotion and may urge them to take action if usage rates appear too low.

Third, a trial for a high-tech service may trigger specific decision dynamics. For one, whereas current research deals with settings in which conversion of triers into paying customers occurs on an *opt-in* basis, users of a temporary service trial typically must explicitly *opt out* if they prefer not to adopt the paid service. As a result, trial users may adopt just to avoid the hassle of opting out. On the other hand, for subscribers to become acquainted with the system, a high-tech service trial usually covers a sufficiently long period; as a consequence, consumers may get used to the trial's zero price and become less willing to adopt the paid service.

In this paper, we examine how free-trial promotions for contractual high-tech consumer services influence adoption. We focus on adoption (rather than repurchase) decisions because they are typically associated with higher risk and lower product knowledge, and are critical to the success of new products (Ho, Li, Park, & Shen, 2012). We present a framework for the trial

mechanisms and develop a model for a consumer's free-trial subscription and paid adoption decisions. Our model accommodates consumers' quality learning from trial usage, WOM, advertising, and direct marketing, in a context in which service quality itself evolves over time. Consumers are forward-looking in that they account for quality changes and anticipate uncertainty reduction due to trial usage. We estimate our model on a unique data set, incorporating customers' trial subscription, adoption, and usage behavior for an interactive digital television (IDTV) service. The interactive TV and video-on-demand market has evolved into a multibillion-dollar industry that spurs the interest of academics (e.g., Nam, Manchanda, & Chintagunta, 2010). Our findings generate important managerial insights. First, we assess the effects of a free trial on consumers' adoption behavior and compare the trial's learning effects with those of advertising and direct marketing. Second, we demonstrate that ill-timed free trials may actually reduce the number of adopters. Third, we shed light on the appropriate timing of free-trial promotions and the role of consumers' usage intensity during the trial.

Below, we first develop our conceptual framework. Next, we elaborate on the study context and available data. We then present the model specification and estimation results, and simulate the performance implications of alternative free-trial scenarios. We conclude with a discussion of academic and managerial implications, limitations, and future research.

2. Conceptual Framework

Figure 1 depicts the processes that govern a consumer's decision to adopt a new high-tech service. In the absence of a free-trial promotion, the consumer's propensity to adopt depends on marketing communication and WOM (Manchanda, Xie, & Youn, 2008): advertising or direct marketing, and contacts with current adopters, enable the consumer to learn about the service's quality (see the solid lines in Figure 1). A temporary free trial also influences adoption, in two

ways (dashed lines in Figure 1): directly, by allowing the consumer to try the service and learn through own usage, and indirectly, because trial subscribers may spread WOM.

Whether a free trial increases the number of adopters depends on the extent to which it generates new subscribers who would not have adopted without the trial. However, as we will show later, the free trial's effectiveness also rests on its ability to avoid losing customers who would have adopted in the absence of a trial promotion. Figure 2 presents the free trial's impact on adoption as the difference between the number of adopters gained and the number of adopters lost. Below, we discuss the components of Figure 2 and link them to the mechanisms in Figure 1.

-- INSERT FIGURES 1 AND 2 ABOUT HERE --

2.1. Adopters Gained in Response to Trial Promotion

In Figure 2, the gray- and black-striped boxes refer to adopters gained after own trial usage and adopters gained merely due to WOM, respectively.

Adopters gained after own trial usage. A free trial may appeal to consumers who otherwise would not subscribe, for several reasons (see Figure 1). First, it allows them to start using the service without paying any setup or periodic fees. Second, in contrast with a paid subscription, which involves a long-term commitment (e.g., a one-year software license), a free trial permits the consumer to end the relationship with the seller sooner, without penalty. Though some consumers may prefer a long-term contract for precommitment or convenience reasons (Wertenbroch, 1998), most consumers value the possibility to reconsider their subscription at a later moment, when they have experienced the service and have become more certain about its quality. That is, the trial provides a so-called option premium (Sriram, Chintagunta, & Manchanda, 2014). Finally, in addition to their purely rational benefits, the trial's free nature and

its option premium may trigger smart-shopper feelings, in which consumers derive pleasure merely from receiving a deal (Bicen & Madhavaram, 2013).

After completing the free trial, consumers may be more likely to become paying adopters (see Figure 1). First, free triers pay no setup fee, because the service is already in place. Second, continuation of the subscription does not require any installation effort; rather, not acting and continuing to use the service implies a reduction of effort, a so-called inertia premium (Su, 2009). Third, and most crucially, triers can learn about the service's quality through own experience, which may improve their perceived service benefits. Learning reduces uncertainty, which has a positive effect on (risk-averse) consumers' appraisal of the service. In addition, learning from trial usage can generate quality beliefs that are more favorable than those based on external quality signals alone. For example, consumers may doubt the credibility of advertising and direct marketing and therefore discount (possibly more than is justified) quality signals obtained via these marketing channels (Mehta, Chen, & Narasimhan, 2008).

Adopters gained due to WOM. A free trial may also generate new adopters in an indirect way (Ghose & Han, 2011). As Figure 1 indicates, triers and adopters-after-trial may engage in WOM communication, such that current nonusers learn more about the service's actual quality (Jain et al., 1995; Manchanda et al., 2008; Nam et al., 2010). Quality signals obtained through WOM help nonusers reduce their quality uncertainty and may lead to more favorable quality beliefs than the less-trusted marketing signals.

2.2. Adopters Lost in Response to Trial Promotion

In the context of high-tech contractual services, free trials can also have *negative* effects on consumers' adoption decisions. That is, some consumers who would have adopted a long-term

contract in the absence of a trial promotion might refrain from adopting after trial usage or due to WOM. In Figure 2, we use the gray- and black-colored boxes to indicate these consumer groups.

Adopters lost after own trial usage. A free trial may attract people who willingly would have adopted the paid service in the absence of the trial. This creates important risks. First, as noted in Figure 1, the free trial may reduce reference prices (Pauwels & Weiss, 2008), such that the fees that come with regular adoption loom larger after a free trial. Second, trial usage enables the consumer to learn and become more certain about the service's quality, but the net impact of such learning is not necessarily positive (Sriram et al., 2014). When the consumer has a negative trial experience, the resulting quality beliefs may be less favorable than those formed (only) on the basis of advertising and direct marketing signals, even when she discounts marketing information. This particularly holds for high-tech services, whose actual performance is evolving and which may suffer from technical malfunctions, especially early in the service life cycle. The consumer may insufficiently account for the temporary nature of these deficiencies and take the bad experience as a signal of the service's long-run quality. As a result, consumers who otherwise would have accepted a regular subscription, may renege after a free trial.

Adopters lost due to WOM. Negative trial experiences may be shared with others, such that also consumers who do not use the trial, learn about the service's deficiencies. These WOM signals make consumers more certain about the service's quality, but also decrease the valence of their quality beliefs. As a result, they may decide not to subscribe to the service, which they would have adopted without the trial.

In summary, the final impact of the free trial rides on the dual mechanism of attracting consumers who otherwise would *not* have subscribed to the service, and minimizing the losses of consumers who otherwise *would* have adopted. A crucial determinant of the relative sizes of

these consumer groups is consumers' learning process in the presence of a trial: what consumers (directly or indirectly) learn from the trial, may differ from what they learn from marketing, and may change over time – aspects that we incorporate in our adoption decision model.

3. Data

3.1. Study Context

We use data from a large West European telecom firm that offers telephone, Internet, and TV services and operates in a single country. The data cover consumers' adoption behavior and the company's marketing efforts for DSL-based interactive digital television (IDTV) during the first two years after the service's launch. In the country and period of study, the company has a market share of 31% in the digital TV market, and is the only provider of IDTV via DSL.² To allow modeling of consumers' adoption decisions, our data also include customers who never subscribed to IDTV and only use the company's other services.

To use the IDTV service, customers must install a DSL modem and set-top box that decodes the digital signal. The service grants access to a wide selection of television channels in digital quality, an electronic program guide, program background information, and video-on-demand—including movies, shows, and newscasts (because the studied provider does not *produce* TV content, user interaction with TV programs was virtually non-existent). New IDTV customers sign a 12-month contract with automatic renewal. They pay a one-time installation fee and a monthly flat rate for the basic channel package and set-top box.

In an attempt to accelerate adoption, the company made intensive use of free-trial promotions: between months 10 and 19 of our two-year observation period, every interested consumer could subscribe to a three-month free trial of the basic IDTV offer, without any

² Its main competitor (with a market share of 40%) and the smaller remaining players in the digital TV market (with a total market share of 29%) operate through satellite or cable.

purchase obligation. After the trial, customers could cancel the contract and return the set-top box to one of the operator's many shops, or else do nothing and become regular paying users. The contract of consumers who did not cancel after a free trial, was automatically converted into a paid one such that the next nine months were considered part of a regular contract.

3.2. Data Set

From our initial set of more than 170,000 consumers, we deleted those who did not technically qualify to receive the IDTV signal (e.g., because they did not have a sufficiently fast DSL Internet connection from the focal company) or for whom relevant information was lacking. We randomly selected 15,000 customers, of whom 10,000 were used for estimation and 5,000 for out-of-sample validation. Because our sample consists of people who were customers of the company prior to the launch of the IDTV service (i.e., they had the required Internet connection), we also can identify nonadopters.

For each consumer, we have information about whether and in which month she subscribed to the free trial or adopted the paid IDTV service. Furthermore, our data set contains information about the four sources through which consumers learn about service quality. Direct marketing and advertising are captured by the monthly number of direct marketing contacts per consumer (phone, e-mail, or regular mail) and monthly per capita advertising expenditures, respectively.³ For usage, we used the monthly number of channel zaps by each subscriber.⁴ Finally, for our WOM measure, we take into account that geographically close neighbors are more likely to interact (Manchanda et al., 2008; Nam et al., 2010) and compute the distance-

³ Because the company's communication management is based on the country's administrative regions, the advertising information is region-specific. Yet, the ad spending differences between regions are negligible. The other marketing instruments also hardly differ across regions: TV channel bundles are largely the same and price and sales promotion strategies are identical.

⁴ This usage measure ensures that the subscriber was actually watching TV. Set-top box operating times, as an alternative, involve both active and passive (i.e., "background noise") use. In addition, IDTV subscribers could leave their set-top box on when their TV set was switched off. Information regarding set-top box operating times is available only for a limited period. During this period, the correlation between zaps and operating times was .67.

weighted number of current IDTV users surrounding each consumer i in month t . Specifically, we compute $\sum_{j \neq i} u_{j(t-1)} / (1 + D_{ij})$, where $u_{j(t-1)}$ is a dummy variable indicating whether consumer j was using the service at the end of the previous month, and D_{ij} is the distance in kilometers between the centroids of the census blocks (i.e., geographic units of 213 households on average) of consumers i and j (e.g., Gauri, Pauler, & Trivedi, 2009).⁵ To reduce the scale of our WOM measure, we divide it by 1,000, without loss of generality. Several other variables served as controls; we discuss them further when we introduce them into the model. In Table 1, we report descriptive statistics for all explanatory variables.

Table 1 also presents some descriptives on free-trial usage and paid adoption behavior. 20.74% of the customers in our calibration set adopted the paid IDTV service at some point during the observation period. Of all customers who took advantage of the free trial, 60.56% became paying IDTV users, compared to only 13.58% of those who did not accept the free trial. These figures suggest that a free trial stimulates paid adoption, but they should be interpreted with caution. First, many adopters-after-trial would have adopted even without the trial promotion. Second, the percentages in Table 1 may mask substantial heterogeneity in consumers' response to free trials. Third, these figures do not address the *dynamic* nature of the trial's effect on adoption. As service quality evolves over time, the outcome of a consumer's learning process—whether through own service usage during the trial or through interactions with current IDTV users—may vary. To cleanly identify the impact of free-trial promotions on a consumer's adoption decision, we need to estimate a model that incorporates these factors.

-- INSERT TABLE 1 ABOUT HERE --

4. Model Development

⁵ We divide by $(1 + D_{ij})$, rather than D_{ij} , to avoid division by zero when consumers i and j belong to the same census block (e.g., Gauri et al., 2009). To increase the reliability of our WOM measure, consumers j involve people both within and outside our calibration data set.

4.1. General Model Structure

We propose a discrete-choice framework in which the consumer, at the end of every month, chooses whether to accept the free trial, adopt the paid offer (possibly after first having used a trial), or not subscribe at all. This framework essentially corresponds to a discrete hazard-rate structure with time-dependent variables (e.g., Manchanda et al., 2008), and accounts for duration dependence in a rich way. First, we disentangle consumers' learning processes preceding adoption: every month, consumers revise their service quality beliefs on the basis of signals from direct marketing, advertising, usage during a trial, and WOM, in the course of which they adjust for possible signal biases (Mehta et al., 2008). Second, consumers account for future quality changes and, when confronted with a free-trial offer, anticipate the learning benefits from the trial experience. Thus, as time passes by and the decision window shifts forward, consumers' assessment of the future service benefits changes.

A few remarks with regard to the availability of consumers' choice alternatives over time are appropriate. First, the free-trial option was only available from months 10 to 19. Second, if a consumer accepts the free trial, the next nontrivial decision comes only at the end of the trial period, when she chooses whether to become a paying user. Third, in the period during which the free trial was available, the company continued offering regular paid subscriptions, typically at a reduced price (e.g., a discount on the installation or monthly fee).⁶ Fourth, in line with our data, consumers who decide not to adopt the paid offer after a trial will not reconsider the service during the remainder of the observation period.⁷ Figure 3 provides an illustration. In the

⁶ Promotional offers could not be accumulated: consumers had to choose the free trial or the discounted paid adoption. Even when the free trial's cost savings outstripped those of the discounted regular offer, some consumers chose the regular offer, most likely because of precommitment or convenience reasons (Wertenbroch, 1998).

⁷ Our data set does not contain a single consumer that opted out after a free trial and later subscribed again to either a trial or the paid offer. This is not surprising: in opt-out-based free trials like those considered here, a consumer's decision not to adopt the paid service tends to be a deliberate choice that she will not easily revise (see Sriram et al., 2014).

beginning, the consumer can only choose between adopting the paid service or not adopting; in this example, the consumer does not adopt. Even during the free-trial campaign (starting in month 10), the consumer does not subscribe until month 14, when she accepts the free-trial offer. The last decision occurs three months later, upon the trial's expiry, when she chooses to adopt the paid offer.

-- INSERT FIGURE 3 ABOUT HERE --

4.2. Discrete-Choice Model

We use a multinomial logit (MNL) structure to model the probabilities that, in a specific month, the consumer subscribes to either the free-trial offer or the regular paid service (possibly after a trial), given that she did not do so before. We write consumer i 's utility in month t for the free-trial (U_{it}^F) and regular offer (U_{it}^R) as follows:

$$(1) \quad U_{it}^F = V_{it}^F + \varepsilon_{it}^F = \beta_i^0 + \beta_i^F - E[\exp(-r_i q_{it}^F)] + \beta_i^{\text{DISC}} \text{DISC}_{it}^F + \beta_i^{\text{UR}} \text{UR}_{it} + X_{it}^F \beta_i^X + \varepsilon_{it}^F$$

$$(2) \quad U_{it}^R = V_{it}^R + \varepsilon_{it}^R = \beta_i^0 - E[\exp(-r_i q_{it}^R)] + \beta_i^{\text{DISC}} \text{DISC}_{it}^R + \beta_i^{\text{FE}} \text{FE}_{it} + X_{it}^R \beta_i^X + \varepsilon_{it}^R.$$

We start by discussing the free trial's utility (Equation 1). First, $-E[\exp(-r_i q_{it}^F)]$ represents the expected service benefits obtained by subscribing to the free trial. In line with previous studies (e.g., Narayanan & Manchanda, 2009), we capture these expected benefits through the expected value of an exponential function of q_{it}^F , which is the consumer's stochastic belief with regard to service quality during the free trial. For positive values of the risk-aversion coefficient r_i , the exponential function (including the negative sign in front of it) is concave, such that the consumer is risk-averse with respect to uncertainty in q_{it}^F . To ensure that r_i is positive, we write $r_i = \exp(r_i^*)$, where r_i^* can take any value. Assuming that q_{it}^F follows the normal distribution $N(\bar{q}_{it}^F, s_{it}^F)$, we can write (see Erdem, Keane, Öncü, & Strebel, 2005):

$$(3) \quad -E[\exp(-r_i q_{it}^F)] = -\exp\left[-r_i \left(\bar{q}_{it}^F - r_i (s_{it}^F)^2 / 2\right)\right],$$

where $-r_i (s_{it}^F)^2 / 2$ is a penalty for quality uncertainty. As we will explain in the next section, the consumer derives her quality belief q_{it}^F by learning from quality signals obtained through advertising, direct marketing, and WOM.

Second, to account for the monetary benefits of the free trial (see Figure 1), we include $DISC_{it}^F$, the sum of all waived installation, activation, and monthly subscription fees over the three trial months. Notice that its coefficient, β_i^{DISC} , is also identified by the occurrence of price promotions for the regular offer (see Equation 2).

Third, UR_{it} is the anticipated utility due to a better-informed adoption decision after the free trial. More precisely, it captures the consumer's anticipated reduction in uncertainty about the service's quality, given her expected usage intensity during the trial (Ching, Erdem, & Keane, 2011). The greater UR_{it} , the more the consumer appreciates the opportunity to make the actual adoption decision only after having tried the service. As such, this component captures the fluctuations in the free trial's option premium. We discuss the derivation of UR_{it} in the next section. Any remaining utility shifts are captured by the coefficient β_i^F . Indeed, smart-shopper feelings may further boost the free trial's utility (see Figure 1).

After a free trial, consumers may consider adopting the paid offer. In Equation 2, two components capture the effect of a consumer's free-trial experience. First, we include $-E[\exp(-r_i q_{it}^R)]$, the expected service benefits of the paid offer, where $q_{it}^R \sim N(\bar{q}_{it}^R, s_{it}^R)$ is the consumer's stochastic belief with regard to service quality during a regular subscription. Unlike q_{it}^F , q_{it}^R is not only based on marketing and WOM signals but also incorporates usage signals if the consumer first subscribed to a trial (see Figure 1). In line with Equation 3, we can write:

$$(4) \quad -E[\exp(-r_i q_{it}^R)] = -\exp \left[-r_i \left(\bar{q}_{it}^R - r_i (s_{it}^R)^2 / 2 \right) \right].$$

Second, we add FE_{it} , a dummy that equals 1 when the consumer has completed a free trial. We thus account for the fixed-effect change in the utility of the paid offer after a trial, due to the absence of a setup fee, the inertia premium, and reference price effects (see Figure 1). The sign and magnitude of the coefficient β_i^{FE} reveal the net impact of these different mechanisms.

Finally, X_{it}^F and X_{it}^R are row vectors of offer-, consumer-, and month-specific control variables. β_i^X is a vector of coefficients. ε_{it}^F and ε_{it}^R are extreme-value distributed error terms, such that we obtain a multinomial logit form for the probability that consumer i subscribes to offer y ($y = F, R$) in month t :

$$(5) \quad p_{it}^y = \frac{z_{it}^y \cdot \exp(v_{it}^y)}{z_{it}^F \cdot \exp(v_{it}^F) + z_{it}^R \cdot \exp(v_{it}^R) + 1}.$$

The probability of non-subscription (N), in turn, is given by:

$$(6) \quad p_{it}^N = \frac{1}{z_{it}^F \cdot \exp(v_{it}^F) + z_{it}^R \cdot \exp(v_{it}^R) + 1}.$$

The indicators z_{it}^F and z_{it}^R serve to “deactivate” an offer that is not available to the consumer. For example, after a trial, the consumer can only choose between the paid offer and the no-purchase option, such that $z_{it}^F = 0$ and $z_{it}^R = 1$ (see Figure 3).

4.3. Consumer Learning About Ultimate Match Quality

At the core of our model is a consumer learning process (Erdem et al., 2005; Narayanan & Manchanda, 2009). Specifically, the consumer tries to infer the *ultimate match quality* θ_i , which is the extent to which the service offers the consumer’s sought-after features *in the absence of any service deficiencies*, that is, after initial technical malfunctions have been resolved. Below, we discuss this learning process in detail. In the next section, we explain how it enables the consumer to derive her beliefs q_{it}^F and q_{it}^R with regard to quality during the trial and regular subscription, and the anticipated uncertainty reduction UR_{it} (see Equations 1 and 2).

Quality information sources. In line with substantial empirical research on learning (e.g., Erdem et al., 2005; Narayanan & Manchanda, 2009), consumers update their beliefs with regard to θ_i in a Bayesian fashion. They take into account quality signals from four sources of information: (1) advertising, (2) direct marketing, (3) own usage of the service during the trial, and (4) WOM (Erdem et al., 2005). $N(\theta_i^A, \sigma^A)$, $N(\theta_i^A, \sigma^D)$, $N(\theta_{it}^U, \sigma^U)$, and $N(\theta_{it}^U, \sigma^W)$ represent the corresponding normal signal distributions. While each of the four sources has its own standard deviation (and thus perceived information precision; see, e.g., Narayanan, Manchanda, & Chintagunta, 2005), advertising and direct marketing have the same mean quality signal θ_i^A (i.e., the advertised match quality) and usage and WOM share the (time-variant) mean quality signal θ_{it}^U (i.e., the real match quality in the current period) (Ghose & Han, 2011).

Signal bias. Importantly, in line with Mehta et al. (2008), these four sources do not necessarily provide *unbiased* signals of the *ultimate* match quality θ_i . First, advertisements and direct marketing contacts often engage in puffery and overstate the service's quality (Jing Xu & Wyer Jr, 2010). In other words, θ_i^A may be greater than θ_i . We capture this by writing θ_i^A as:

$$(7) \quad \theta_i^A = \theta_i + \delta,$$

where δ is the bias in the advertising and direct marketing signals.

Second, especially in the early stages of the service's life cycle, the service may still suffer from technical problems, such that the actually delivered quality, θ_{it}^U , may be below the ultimate match quality, θ_i . We thus write θ_{it}^U as follows:

$$(8) \quad \theta_{it}^U = \theta_i + \kappa \cdot \text{intervent}_t,$$

where $\kappa \cdot \text{intervent}_t$ represents the deviation from the ultimate quality, i.e. the "bias" in the usage and WOM signals. The variable intervent_t refers to the number of onsite repair interventions by a field technician in month t , expressed as a percentage of the total number of customers in that

month.⁸ Including intervent_t helps us to identify the fluctuations in θ_{it}^U . It captures the service's technical evolution and tends to decrease over time, varying between 6.0% and 1.6%. We expect κ to be negative: a high value for intervent_t is an indication of a low current match quality.

Signal correction. Consistent with Mehta et al. (2008), the consumer is aware of possible biases and holds beliefs about the degree of distortion, which enable her to correct the quality signals. As in Mehta et al. (2008), we assume that these beliefs about signal bias, themselves, are not subject to learning. The consumer's belief about the advertising and direct-marketing bias, δ , is represented by d , which follows the normal distribution $N(\bar{d}, s^d)$. In addition, the consumer holds beliefs about $\kappa \cdot \text{intervent}_t$, i.e., the extent to which the real match quality θ_{it}^U deviates from the ultimate match quality θ_i in a given month t . Because the consumer does not observe intervent_t , she assumes the deviation to decrease linearly over time. Specifically, the consumer believes the difference between real and ultimate match quality equals

$$(9) \quad \begin{array}{ll} \kappa \cdot (T - t) & \text{for } t \leq T, \text{ and} \\ 0 & \text{for } t > T. \end{array}$$

The parameter k captures the rate at which, according to the consumer, the technical deficiencies die down; k is stochastic and follows the normal distribution $N(\bar{k}, s^k)$, where \bar{k} is expected to be negative. T represents the number of months it takes for the system to start functioning optimally (i.e. without deficiencies). To keep estimation tractable, we assume that the consumer's belief about T is deterministic. Note that Expression 9 can also be written as $k \cdot \max(0, T - t)$.

Based on her beliefs about the amount of bias, the consumer corrects the incoming quality signals in an attempt to obtain an unbiased estimate of θ_i . Specifically, when the consumer receives some advertising (or direct marketing) signal θ^* , she will use the corrected

⁸ intervent_t does not include customers' requests for *remote* assistance. Typically, such requests can be attributed to customers' initial inability to operate the system and therefore are themselves subject to consumer learning. Our measure is free of such learning mechanisms and thus offers a good indicator of the service's objective performance.

signal $\theta^* - d$. Similarly, when she receives some usage (or WOM) signal $\theta^\#$ in month t , she will rely on the corrected signal $\theta^\# - k \cdot \max(0, T - t)$. Given Equations 7 and 8, the means of the corrected signal distributions are $\theta_i + \delta - \bar{d}$ for advertising and direct marketing, and $\theta_i + \kappa \cdot \text{intervent}_t - \bar{k} \cdot \max(0, T - t)$ for usage and WOM.

Bayesian updating. Let \mathcal{G}_{it}^A , \mathcal{G}_{it}^D , \mathcal{G}_{it}^U , and \mathcal{G}_{it}^W be the sums of the corrected advertising, direct marketing, usage, and WOM signals that the consumer obtains in month t . If we condition on d and k (a constraint that we will later relax again), these signal sums have the following normal distributions:

$$(10) \quad \begin{aligned} \mathcal{G}_{it}^A | d &\sim N\left(n_{it}^A \cdot (\theta_i + \delta - d), \sqrt{n_{it}^A} \cdot \sigma^A\right), \\ \mathcal{G}_{it}^D | d &\sim N\left(n_{it}^D \cdot (\theta_i + \delta - d), \sqrt{n_{it}^D} \cdot \sigma^D\right), \\ \mathcal{G}_{it}^U | k &\sim N\left(n_{it}^U \cdot [\theta_i + \kappa \cdot \text{intervent}_t - k \cdot \max(0, T - t)], \sqrt{n_{it}^U} \cdot \sigma^U\right), \text{ and} \\ \mathcal{G}_{it}^W | k &\sim N\left(n_{it}^W \cdot [\theta_i + \kappa \cdot \text{intervent}_t - k \cdot \max(0, T - t)], \sqrt{n_{it}^W} \cdot \sigma^W\right), \end{aligned}$$

where n_{it}^A , n_{it}^D , n_{it}^U , and n_{it}^W refer to the numbers of, respectively, advertising, direct marketing, usage, and WOM signals in month t . The above expressions hold true because the original, uncorrected, signals follow the normal distributions $N(\theta_i + \delta, \sigma^A)$, $N(\theta_i + \delta, \sigma^D)$, $N(\theta_i + \kappa \cdot \text{intervent}_t, \sigma^U)$, and $N(\theta_i + \kappa \cdot \text{intervent}_t, \sigma^W)$. Expressions 10 show that, conditionally on d and k , the corrected signal sums are independent. If we assume that consumers' initial belief about θ_i , q_0 , is normally distributed, then the posterior conditional belief at any point in time is normal too (Narayanan & Manchanda, 2009): $q_{it} | d, k \sim N(\bar{q}_{it}, s_{it})$, with

$$(11) \quad \bar{q}_{it} = s_{it}^2 \left(\frac{\bar{q}_{i(t-1)}}{s_{i(t-1)}^2} + \frac{\mathcal{G}_{it}^A}{(\sigma^A)^2} + \frac{\mathcal{G}_{it}^D}{(\sigma^D)^2} + \frac{\mathcal{G}_{it}^U}{(\sigma^U)^2} + \frac{\mathcal{G}_{it}^W}{(\sigma^W)^2} \right) \text{ and}$$

$$(12) \quad s_{it} = \left(\frac{1}{s_{i(t-1)}^2} + \frac{n_{it}^A}{(\sigma^A)^2} + \frac{n_{it}^D}{(\sigma^D)^2} + \frac{n_{it}^U}{(\sigma^U)^2} + \frac{n_{it}^W}{(\sigma^W)^2} \right)^{-1/2}.$$

Notice that the consumer can use $q_{it}|d, k$ to derive a conditional belief with regard to the real match quality in some future period $t + \tau$, that is, with regard to $\theta_{i(t+\tau)}^U$. Let $q_{it}^{t+\tau}|d, k$ refer to this conditional belief. Given the consumer's belief about the difference between current and ultimate match quality (see Expression 9), we obtain $q_{it}^{t+\tau}|d, k = (q_{it}|d, k) + k \cdot \max(0, T - t - \tau)$, which is normally distributed.

4.4. Derivation of q_{it}^F , q_{it}^R , and UR_{it}

In our model, consumers are forward-looking in two ways. First, they account for the fact that match quality θ_{it}^U is evolving. Specifically, in their adoption and trial subscription decisions, they use their beliefs about the quality changes to derive the beliefs q_{it}^F and q_{it}^R with regard to the *average* service quality during the free trial and regular subscription, respectively (see Equations 1 and 2). Second, when evaluating the trial offer, consumers anticipate the uncertainty reduction UR_{it} as a result of learning during the trial. Below, we discuss both elements.

Derivation of q_{it}^F and q_{it}^R . In Web Appendix 1, we obtain the unconditional normal distributions of q_{it}^F and q_{it}^R by averaging the conditional match quality beliefs $q_{it}^{t+\tau}|d, k$ across the months $(t + 1, t + 2, \dots)$ of the free trial and regular subscription, respectively, and then deriving the unconditional mean and standard deviation. For the free trial, we find:

$$(13) \quad \bar{q}_{it}^F = \sum_{\tau=1}^{L^F} E(q_{it}^{t+\tau}|d = \bar{d}, k = \bar{k})/L^F \text{ and}$$

$$(14) \quad s_{it}^F = \left[s_{it}^2 + G_{it}^2 \cdot (s^d)^2 + H_{it}^2 \cdot (s^k)^2 \right]^{1/2}.$$

The expressions for \bar{q}_{it}^R and s_{it}^R are analogous. $q_{it}^{t+\tau}|d = \bar{d}, k = \bar{k}$ refers to $q_{it}^{t+\tau}|d, k$ evaluated in \bar{d} and \bar{k} , L^F represents the length in months of the free trial, and G_{it} and H_{it} are weights that depend, among others, on the numbers of past quality signals and signal variances (see Web Appendix 1 for the full expressions). Equation 14 shows that the uncertainty about the average match quality

during the trial, $(s_{it}^F)^2$, depends not only on the uncertainty about the ultimate match quality, s_{it}^2 , but also on the uncertainty about signal bias, $(s^d)^2$ and $(s^k)^2$ (see Mehta et al., 2008).

Derivation of UR_{it} . Following Ching et al. (2011), we model the expected uncertainty reduction as the anticipated decrease in the variance of the consumer's quality belief as a result of service usage.⁹ More precisely, the consumer evaluates the extent to which usage during the trial will affect $(s_{i(t+3)}^R)^2$, which is the uncertainty three months from now about the average match quality during a regular subscription starting in $t + 4$. Thus, the consumer assesses the difference between, on the one hand, the anticipated uncertainty in the absence of any usage signals and, on the other hand, the anticipated uncertainty given the consumer's likely usage pattern during the trial. As we explain in Web Appendix 1, to derive this uncertainty reduction, the consumer needs the expected number of advertising, direct marketing, usage, and WOM signals during the next three months, for which she relies on her signal history and consumer characteristics. Finally, to cleanly separate the anticipated uncertainty reduction over time from cross-sectional differences between consumers (which are already captured by the heterogeneous utility intercept), we center each customer's series of uncertainty reduction values around her first-period value. As a result, UR_{it} only captures within-consumer cross-time variation in anticipated uncertainty reduction.

⁹ In line with Ching et al. (2011) and Ching, Erdem, and Keane (2013), we thus use a "reduced form" approach to capture future benefits, rather than solve the consumer's dynamic programming problem. Not only is the latter approach more complicated and time-consuming, identification may become prohibitive if the value outcomes of the decisions are not directly observed, which is the case in our model, since we only observe choices, not the underlying utilities (Ching et al., 2011). Moreover, the only learning effects that a consumer can truly anticipate are those on belief *uncertainty*, and not those on the belief's *expected value*, for which she would need the exact values of the future quality signals (Ching et al., 2011). Hence, one can express the future learning benefits of making a certain choice (in our case, accepting the free trial) as a function of the reduction in uncertainty following that choice compared to the no-purchase option. To test the accuracy of our reduced form, we run a simulation in which we compute the anticipated uncertainty reduction due to trial and compare it with the anticipated benefit of trial in a more structural model. The approach and results are discussed in Web Appendix 2.

Notice that the role of UR_{it} differs fundamentally from that of s_{it}^F which also appears in the free trial's utility function (see Equations 1 and 3). s_{it}^F captures uncertainty with regard to the average match quality during the free trial, given the quality signals collected up till month t . In contrast, UR_{it} measures uncertainty *reduction*, pertains to the average match quality in the months *after* the trial, and accounts not only for the so-far collected quality signals but also for the *anticipated* signals in the next three months.

In summary, we build a model in which consumers learn about ultimate match quality (i.e., match quality in the absence of technical deficiencies) from advertising, direct marketing, WOM, and, in case they subscribe to the trial, own service usage. Consumers expect the advertising and direct marketing signals to overstate the service's quality and therefore adjust them downward, relying on their beliefs about puffery in marketing communication. Furthermore, consumers know that the WOM and usage signals may *understate* the ultimate quality due to temporary technological deficiencies, and correct these signals based on how much they believe the actual quality is still below the ultimate one. The derived belief about ultimate match quality together with the belief about the speed of technological evolution, enable consumers to make inferences about the average quality during a future (paid or free-trial) contract, which informs their subscription decisions. In addition, when considering a free trial, consumers account for the extent to which the trial will help reduce the uncertainty about average quality.

4.5. Control Variables

Several other factors, represented by X_{it}^F and X_{it}^R in Equations 1 and 2 respectively, directly influence the utilities of the trial and paid offer. First, we include advertising and direct marketing in both utility equations to account for these variables' persuasive role, in addition to their informative signaling effects on consumers' quality beliefs (Narayanan et al., 2005).

Second, we account for the fact that, near the end of our observation period (after the free-trial campaign), a *paid* three-month trial option became available, allowing consumers to experience the paid service without signing a long-term contract. Although our model already captures learning through usage and WOM, two additional variables are necessary to accommodate a paid trial's effects. Specifically, in the trial's utility function, we include a dummy that equals 1 when the trial was not free of charge; this dummy's coefficient replaces β_i^F in Equation 1. We thus allow for any difference in utility between free and paid trials (e.g., due to limited smart-shopper feelings), on top of the utility shift captured by the discount variable $DISC_{it}^F$. Similarly, in the utility function of the regular contract, we incorporate a dummy (comparable to FE_{it}) that equals 1 when the consumer has completed a paid trial; the fixed feedback effect of a paid trial may be higher than that of a free trial because a paid trial is unlikely to decrease reference prices. We only add these variables as controls; free trials remain our focus.¹⁰ Third, in the utilities of the regular and trial offer, we include monthly advertising expenditures (euro per capita) of the main competitor and a linear trend to account for any remaining changes in subscription probability.

4.6. Consumer Heterogeneity

To account for latent consumer heterogeneity, the coefficients β_i^F , β_i^{DISC} , β_i^{UR} , β_i^{FE} , and β_i^X , the transformed risk-aversion parameter r_i^* ($r_i^* = \ln(r_i)$), and the ultimate match quality θ_i (and, as a result, also θ_i^A and θ_{it}^U , see Expressions 7 and 8) follow normal mixing distributions with constant population-level means and standard deviations. The intercept β_i^0 is also normally distributed but its mean is a function of several consumer characteristics:

$$(15) \quad E(\beta_i^0) = \alpha^0 + Z_i \alpha^Z,$$

¹⁰ The paid trial was available only toward the end of our observation period, so the data set contains very few customers who actually completed such a trial.

where α^0 (α^Z) is a (column vector of) coefficient(s) to be estimated, and Z_i is a row vector of (mean-centered) consumer characteristics, including the consumer's age and relationship length (i.e., the duration of the consumer's DSL subscription) at the start of the observation period, average annual income in the consumer's census block, and family size (see also Table 1).

Several of these potential sources of (latent or explained) consumer heterogeneity are shared by the free trial and paid service (see Equations 1 and 2), such that our model allows for correlation between the offers' utilities (e.g., Train, 2009). This is important because consumers who subscribe to the trial offer may have a greater propensity for paid adoption already prior to the trial, e.g., due to their intrinsic interest in the service. By accounting for such commonalities in the utility of both offers, we avoid bias in the estimated effect of trial usage on paid adoption.

4.7. Identification and Estimation

For identification purposes, we set consumers' initial quality uncertainty s_0 to .01 (e.g., Narayanan et al., 2005). Furthermore, δ and \bar{d} cannot be identified individually: marketing bias and consumers' subsequent correction always co-occur (see, e.g., Expressions 10) and neither process can be observed directly (Mehta et al., 2008). We therefore define $d^* = \delta - d$, capturing the amount of over- or undercorrection, and estimate the parameters of the corresponding distribution $N(\bar{d}^*, s^d)$. This identification problem does not arise for the usage and WOM signals, because the bias $\kappa \cdot \text{intervent}_t$ and correction $k \cdot \max(0, T - t)$ are informed by two separate, observable processes: the fluctuation in repair interventions and the passage of time t , respectively. Finally, we assume that T , consumers' expected number of months for the system to become void of technical deficiencies, is in line with the actual pattern of repairs. We thus plot intervent_t against time, and set T equal to the point where the fitted regression line crosses the

time axis (i.e., 29 months after launch). Sensitivity analysis shows that our results are largely insensitive to changes in T . Web Appendix 3 gives a more detailed discussion on identification.

From the researcher's perspective, the parameters $\beta_i^0, \beta_i^F, \beta_i^{DISC}, \beta_i^{UR}, \beta_i^{FE}, \beta_i^X, r_i^*, \theta_i, \eta_{it}^A, \eta_{it}^D, \eta_{it}^U$, and η_{it}^W are random variables that must be integrated out for each consumer. We therefore estimate our model with simulated maximum likelihood, using 100 Halton draws from the parameters' distributions. The log likelihood expression can be found in Web Appendix 4.

5. Results

To establish validity, we first compare our full model (FM) with four benchmark specifications: (1) a model without quality learning but with a dummy to capture the fixed effect of a preceding free trial (BM1); (2) a learning model in which consumers do not look forward, and the quality signals of all sources are unbiased with the same time-invariant mean (BM2); (3) a model that adds anticipated uncertainty reduction to BM2 (BM3); and (4) a model that adds marketing bias (correction) to BM3 (BM4). Table 2 reports the models' log likelihood, BIC, and AIC values, as well as the root mean square error (RMSE) of the predicted monthly hazard rates in both the calibration and the holdout sample. Except for BIC (which favors BM4), the fit measures support our full model and underscore the importance of accommodating learning, signal bias (correction), and forward-looking. As another validity check, we track the observed trial subscription and post-trial adoption numbers over time, against the predictions of our full model (for which we use consumer-specific parameters, based on a Bayesian update of the population-level parameters, see Train, 2009), and find that the model performs quite well (see Figure 4).

Table 3 reports the parameter estimates for our full model. The coefficients of the control variables have face validity. For example, large households have a higher probability to subscribe ($p < .01$)—which makes intuitive sense, given the nature of the service— and the

positive mean coefficients of own advertising ($p < .10$) and direct marketing ($p < .01$) point to a persuasive role of marketing communication. Below, we first discuss the utility parameters that capture the free-trial mechanisms shown in Figure 1. We then zoom in on the learning process underlying the changes in perceived service benefits. For heterogeneous parameters, we focus on the population means: for most parameters, latent consumer heterogeneity is relatively limited.

-- INSERT TABLE 2, TABLE 3, AND FIGURE 4 ABOUT HERE --

5.1. Adoption Decision Mechanisms in the Presence of a Free Trial

Table 3 shows that, as expected, the discount effect is positive ($p < .01$). Hence, monetary benefits influence service adoption and, in the case of a free trial, people are attracted by the opportunity to use the service for free. The positive coefficient of uncertainty reduction ($p < .10$) indicates that consumers value the opportunity to become more certain about service quality before committing to paid adoption. Interestingly, the free-trial constant β_i^F is also positive, and strongly significant ($p < .01$). Thus, free trials have appeal beyond their anticipated monetary and informative benefits and seem to trigger smart-shopper feelings.¹¹ Together, these effects show that the free-trial offer convinces otherwise uninterested consumers to start using the service.

However, it remains to be seen to what extent trial users are converted into paying adopters. First, the fixed-effect coefficient β_i^{FE} is positive and significant ($p < .01$): for free-trial subscribers, any negative reference price effect thus tends to be more than compensated by the inertia premium and the absence of setup fees. Second, this effect may be enhanced or tempered by changes in perceived service benefits as a result of quality learning during the trial. Indeed, the risk-aversion parameter r_i is significant ($p < .01$), such that changes in consumers' quality beliefs due to trial usage affect the attractiveness of the paid contract (see Equation 4). Notice

¹¹ The paid-trial constant is lower than the free-trial constant ($p < .01$), indicating that, even after controlling for the discount, the paid trial appears less attractive, most likely due to limited smart-shopper feelings. The fixed effect *after* paid trials is higher, probably because of the absence of reference price effects.

that r_i also captures changes in quality beliefs due to learning from WOM. Thus, even if a consumer does not subscribe to the trial, she is indirectly affected by the trial through WOM.

5.2. Learning Process

The results in Table 3 also throw light on consumers' learning process. The initial mean belief about ultimate match quality, $\bar{q}_0 = -.594$, is significantly below the average consumer's *true* ultimate match quality $\bar{\theta}_1 = 5.195$ ($p < .01$). Based on Equation 12 and the signals' standard deviations in Table 3, we can determine the effectiveness of the various information sources at updating the beliefs and reducing uncertainty. For instance, a 10% drop in the uncertainty about ultimate match quality s_{it} (i.e., from its initial value of .01 to .009) requires 96.687 zaps, compared to 2.649 advertising signals, 1.195 direct marketing contacts, and 3.164 WOM signals. Viewed against the average monthly levels of usage (157.154 zaps), advertising (€202 per capita), direct marketing (.206 contacts), and WOM (.342), this suggests that consumers' own experiences are by far the most informative quality signals. In line with insights from cognitive psychology, firm communications are less informative than own usage (Hoch, 2002). The WOM signals' relatively low informativeness is also expected because consumers may not actually pick up or fully observe the information signals from prior subscribers (Ghose & Han, 2011). Figure 5 demonstrates that usage can substantially decrease uncertainty about ultimate match quality. The solid line presents s_{it} if the consumer would have completed a three-month free trial in months 4, 5, etc., while the dotted curve shows the evolution of s_{it} absent any service usage.¹² For example, for a free trial ending in month 8, s_{it} equals .006, which is much (i.e., 25%) lower than the corresponding value without preceding trial, namely .008. Notice, however, that for trials later in

¹² In our computations, we use the overall average number of advertising, direct marketing, and usage signals, and take the average consumer's pattern of WOM signals.

time, this uncertainty reduction decreases such that the two lines tend to converge: the extra information value of a trial decreases as consumers cumulate quality signals from other sources.

Table 3 shows that $\bar{d}^* = \delta - \bar{d}$, the difference between actual marketing bias and consumers' mean bias correction, is negative and significant ($p < .01$). Hence, consumers not only discount, but overcorrect the information from advertising and direct marketing. Even so, they trust their signal correction because the standard deviation of their belief about marketing bias is insignificant ($p > .10$). The coefficient κ , which relates the actual match quality θ_{it}^U to the percentage of onsite repair interventions (see Equation 8), is significant ($p < .01$) and has the expected negative sign. The estimate of \bar{k} , which captures consumers' mean *belief* about the monthly quality improvements, is negative and significant ($p < .05$). Hence, consumers realize that the quality signaled by usage and WOM may be substantially below the ultimate match quality, especially shortly after service launch (see Equation 9). Still, their belief is subject to considerable uncertainty because s^k is relatively large ($p < .01$).

So, even though consumers correct incoming signals in an attempt to learn about the ultimate match quality, the corrected signals do not necessarily have a unique, constant signal mean (see Mehta et al., 2008). This becomes clear in Figure 6, which portrays the corrected mean usage and WOM signal values ($\theta_i + \kappa \cdot \text{intervent}_t - \bar{k} \cdot \max(0, T - t)$) and the corrected mean advertising and direct marketing signals ($\theta_i + \delta - \bar{d}$) for the average consumer. The graph shows that, in the early months (especially before month 10), usage and WOM tend to generate less favorable corrected signals than do advertising and direct marketing, while the opposite holds in later periods (especially after month 17). While learning through usage and WOM, consumers insufficiently correct for deficiencies experienced in the early stages and tend to take bad usage experiences as a signal of the service's ultimate match quality. As a result, whether

and how much the free trial enhances or reduces perceived service benefits will strongly depend on the timing of the trial and the amount of learning. We address this more extensively in the next section. Web Appendix 5 provides further checks on the validity of the estimation results.

-- INSERT FIGURES 5 AND 6 ABOUT HERE --

6. Simulations

Though the parameter estimates give a flavor of the mechanisms driving the free-trial impact, the ultimate outcome depends on the number of customers gained and lost, whether through own trial usage or WOM (see Figure 2). In this section, we assess the net impact of free trials and use simulations to explore how this net impact changes with promotion timing and usage intensity.

For each scenario, we simulate 100 decision trajectories per consumer (i.e., we draw signals and trial acceptance and adoption outcomes) during the first two years after launch, while accounting for consumer interdependencies due to WOM. For the random coefficients, we use consumer-specific posterior values (Train, 2009). To avoid confounding promotion effects, we assume away any discounts on the paid offer.¹³ We obtain the expected monthly number of paying service adopters by averaging across the 100 draws and adding up across consumers.

6.1. Effect of a Free-Trial Promotion over Time

To illustrate the dynamic effects of a free trial, Figure 7 displays paid adoption over time for the baseline scenario in which no trial promotion takes place (dashed line), and for the scenario in which a three-month free-trial offer is available to all customers between months 9 and 11 (solid line). The figure shows that during these months, the number of paying adopters decreases

¹³ The values for the other background variables are set as follows. Advertising, direct marketing, and usage rates equal consumer-level averages. For consumers who never subscribed to the service, we compute the usage rates as an exponential function of their sociodemographics, calibrated with the data of actual users. Notice that, in reality, firms may adjust their direct marketing and advertising decisions as a function of the timing of the trial, such that our simulation results could be seen as lower bounds of the true results. Because the percentage of onsite repair interventions may be subject to learning on the part of the firm, we compute intervention as an exponential function of time and (simulated) number of subscribers, calibrated on the data.

because some consumers who would have become paying users in the absence of a trial, now first take advantage of the free offer. Over the next three months, several consumers who have completed the trial or hear about the service become paying users; some of these would not have adopted without a trial promotion. In month 15, when the last trial subscriptions have expired, the number of paying adopters drops sharply below the baseline. While this may be partly due to acceleration (consumers having adopted earlier), some consumers who would adopt the paid offer in the baseline scenario are truly lost. They subscribe to the trial but, being disappointed about the service's quality, defect for good. Or, they decide not to adopt because they receive negative WOM. The cumulative number of paying adopters by the end of month 24 is thus slightly lower than in the baseline scenario (495 versus 503). Further exploration reveals that the difference between the baseline and free-trial scenario results from, on the one hand, attracting 90 adopters who would not have adopted the service without the free trial and, on the other hand, losing 98 customers who would have adopted without the trial. The question then becomes: how to reduce the number of customers lost and/or increase the number of customers gained?

-- INSERT FIGURE 7 ABOUT HERE --

6.2. Impact of Promotion Timing and Usage Intensity

Using simulations, we show how appropriate timing and increased use rates can improve the performance of a free trial. We fix the length of the period during which the consumer can subscribe to the trial at three months and, like before, use a simulation horizon of two years after launch. Other promotion durations and time horizons lead to similar patterns. Our focal performance metric is the total number of adopters at the end of the planning horizon, but when relevant, we also report the total number of months of paid service usage across adopters. The

latter measure takes the *moment* of paid adoption into account and attaches more weight to early adopters, who start generating revenue earlier in time (Bawa & Shoemaker, 2004).

Impact of timing. Figure 8 reveals how a change in the starting month of the trial campaign, influences the number of adopters.¹⁴ The lower part of the graph (Y-axis on the right) shows that early free-trial promotions (i.e., before month 10) lead to more customers lost than gained, while in later periods, the number of customers gained prevails. The upper part of the graph (Y-axis on left) indicates that early trials thus substantially decrease the total number of adopters relative to the baseline (by up to 40%), whereas later free trials increase adoption. Two phenomena favor trials later in the decision window. First, in later periods, a free trial leads to more positive impressions than does advertising or direct marketing, because the technical service quality and corrected usage and WOM signals improve over time (see Figure 6). Second, people who do not adopt after trial are lost for good. As a result, an *early* free trial may kill many consumers who otherwise would have adopted the paid system in one of the remaining months. Trials at later points in time may also kill consumers, but the chance that these consumers would otherwise have adopted in one of the few remaining months is relatively small.

However, as shown in the lower part of Figure 8, the number of adopters gained levels off, and even slightly drops, at a certain point. First, in later months, subscribing to the free trial loses appeal because the expected uncertainty reduction (the option premium) decreases as the number of already collected advertising, direct marketing, and WOM signals grows. Second, consumers who do subscribe to the trial may hardly update their quality beliefs, which may have reached convergence already. For the same reason, any additional WOM signals triggered by the free-trial campaign may leave consumers' beliefs largely unaffected. As a consequence, there is

¹⁴ Notice that month 19 is the last starting month for which we can assess the impact, because the trial is available for three months (e.g., 19, 20, 21), and trial subscribers make a final adoption decision only after three months of service usage (e.g., month 24 if they accept the trial in month 21).

a point in time (month 17 in this case) after which further delay in the trial campaign no longer enhances the total number of adopters. To assess the significance of the differences between the baseline and free-trial curve, we rerun the simulation for each of 100 draws from the parameter estimates' distributions. We derive the 95%-confidence intervals of the differences between the two curves and find that, except in month 9, all differences are significant.

Figure 9 shows the effect of timing on the number of paid adoption months. Analogous to Figure 8, the lower graph presents the number of paid adoption months lost and gained as a function of the starting month of the free-trial campaign. Paid months are lost when consumers do not adopt in the presence of a trial promotion but would have adopted otherwise, yet also when the free trial reaches consumers who would have been willing to immediately adopt the paid service (subsidization). In turn, paid months are gained when people who would not have subscribed otherwise adopt in the free-trial scenario, but also when the trial makes consumers adopt earlier than in the baseline scenario (acceleration). As the lower part of Figure 9 shows, the number of months gained starts to drop for trials offered in month 14 already. Although offering a free trial later in time may still increase the number of adopters (see Figure 8), it also reduces the number of paid months per attracted customer. That is, a trial later in time benefits less from accelerated payments while it continues to subsidize consumers that would have adopted anyway. The upper part of the graph shows that if the objective is to maximize the total number of adoption months, the trial campaign should be scheduled earlier (in month 14 in this case) than when the focus is on the number of adopters.¹⁵ Like in Figure 8, all differences between the baseline and free-trial performance are significant at the 95% level, except in month 9.

-- INSERT FIGURES 8, 9 AND 10 ABOUT HERE --

¹⁵ Recent work by Datta, Foubert, and Van Heerde (2015) shows that customers attracted through free trials may exhibit higher churn rates. Therefore, we reran the simulation using the average retention rates reported in Datta et al. (2015), namely .93 for free-trial and .96 for regular customers. The resulting graph looks similar to Figure 9.

Impact of usage intensity. To further improve free-trial performance, the firm can boost consumers' use rates during the trial, e.g., by temporarily giving them access to extra channels, or granting them credit for a video-on-demand service. Figure 10, Panel a, shows how changes in consumers' usage levels affect the total number of adopters, for a campaign starting in month 17 (the optimal starting month according to Figure 8). As the graph shows, increased usage during the trial period enhances conversion into paid adopters, while lower usage rates (e.g., due to the temporary unavailability of certain channels) trigger the opposite effect. The higher the use rates, the more consumers learn from the trial. The multitude of experiences reduces uncertainty and, for trials starting in month 17 and later, also improves consumers' mean belief about the ultimate service quality (as from month 17, the average corrected usage signal always exceeds the average corrected advertising or direct marketing signal; see Figure 6). In contrast, in earlier periods, learning through usage is not necessarily beneficial: although it still reduces uncertainty, the resulting quality signals may lower consumers' mean quality beliefs. Hence, in those early stages, boosting consumers' usage intensity is not without risk. This is illustrated in Figure 10, Panel b, which depicts the relationship between usage intensity during the trial and paid adoption for a free-trial campaign starting in month 5 (the least favorable starting month according to Figure 8). It shows that for trials offered early on, if anything, increased usage *lowers* adoption. All differences between the baseline and free-trial performance are significant at the 95% level. The insights do not change when we measure performance in terms of paid adoption months.

7. Conclusion

Free trials have become common practice for a wide range of consumer services. Especially in the context of high-tech service contracts, marked by quality uncertainty and substantial setup costs, the opportunity of trying the service for free may appeal to a large number of potential

customers. However, whether this also leads to more paying adopters, and how performance depends on the timing of the offer and usage intensity during the trial, has remained unclear. We address these questions by studying the effects of free trials for an interactive digital TV service.

Our research contributes to a better understanding of the mechanisms that drive the performance of a trial promotion. We show that free trials enhance consumers' propensity to start using the service. The waiver of subscription and setup fees, along with the flexibility to cancel in case of dissatisfaction, entice consumers to take action and subscribe. These trial users may convert into paying adopters because, compared with adoption without prior trial, continuing the subscription does not entail any setup fee and implies lower decision and no installation efforts (the trial's inertia premium). At the same time, the free trial may reduce consumers' internal reference prices, such that the regular subscription fee appears more prohibitive. Our results suggest that the positive inertia effects outweigh the negative reference price effects.

However, the actual conversion rate also critically depends on the changes in perceived service benefits due to learning during the trial. We find that free trials are much more effective at conveying information about the service than are advertising or direct communication: actual usage reduces uncertainty and influences consumers' quality beliefs at a faster pace. Moreover, customers attracted through the trial may "spread the word" and thus also allow current nonusers to learn about the service. As a result, if the service quality is high, the free trial enhances the *number* of adopters but also accelerates many consumers' *moment* of adoption. However, if the service is not (yet) up to standards, the trial offer may be detrimental. Even when consumers realize that quality evolves over time, they may not sufficiently account for the temporary nature of any service deficiencies and take the bad usage experience as a signal of the service's ultimate

quality. In such case, the trial may alienate consumers and trigger adverse WOM effects – driving away customers who would have adopted now or later.

Fortunately, managers can avoid this through appropriate promotion timing. The trial promotion should take place after the service has been tried and tested in the field, to ensure a better trial experience and lower the risk of prematurely killing customers. Moreover, for such well-timed promotions, we show that stimulating usage intensity during the trial (for instance, by granting access to extra channels) may further enhance the subsequent conversion into paid adoption: intensified usage enables triers to benefit more from uncertainty reduction and update their beliefs at a faster pace.

Our research has several limitations that offer avenues for further research. First, additional work could enrich our findings through laboratory-based choice experiments. By including process measures, such experiments can help untangle some free-trial mechanisms (e.g., the inertia premium). Second, we employ a “lost-for-good” framework; a consumer who defects after trial does not return because opting out is an effortful and deliberate action. In settings in which consumers’ choices are not sticky and dropouts can reenter the system, the optimal timing of free trials likely moves to earlier periods, because defection after trial, which especially occurs in the earlier stages, becomes less detrimental. Third, we do not consider the possibility that consumers, once they have subscribed to the service, may use service options for extra charges, such as video-on-demand. This might affect the profitability of free trials. Finally, our model could be extended in several ways. In certain settings, usage intensity may vary with consumers’ quality beliefs. Also, the company’s direct marketing and advertising decisions may be related to the timing of the free-trial campaign. These interdependencies could be modeled and incorporated in policy simulations – extensions that we leave for future research.

References

- Bawa, K., & Shoemaker, R. (2004). The Effects of Free Sample Promotions on Incremental Brand Sales. *Marketing Science*, 23(3), 345-363.
- Bicen, P., & Madhavaram, S. (2013). Reserach on Smart Shopper Feelings. *Journal of Marketing Theory & Practice*, 21(2), 221-234.
- Chatterjee, S., Carroll, E., & Spencer, D. (2009). *Netflix*. Retrieved from London (ON):
- Ching, A. T., Erdem, T., & Keane, M. P. (2011). *Learning Models: An Assessment of Progress, Challenges and New Developments*. Working paper. Rotman School of Management, University of Toronto. Toronto.
- Ching, A. T., Erdem, T., & Keane, M. P. (2013). Learning Models: An Assessment of Progress, Challenges, and New Developments. *Marketing Science*, 32(6), 913-938.
- Datta, H., Foubert, B., & Van Heerde, H. J. (2015). The Challenge of Retaining Customers Acquired with Free Trials. *Journal of Marketing Research*, 52(2), 217-234.
- Erdem, T., Keane, M. P., Öncü, T. S., & Strebel, J. (2005). Learning About Computers: An Analysis of Information Search and Technology Choice. *Quantitative Marketing and Economics*, 3(3), 207-246.
- Gauri, D. K., Pauler, J. G., & Trivedi, M. (2009). Benchmarking Performance in Retail Chains: An Integrated Approach. *Marketing Science*, 28(3), 502-515.
- Ghose, A., & Han, S. P. (2011). *A Dynamic Structural Model of User Learning on the Mobile Internet*. NET Institute Working Paper. Retrieved from <http://ssrn.com/abstract=1485049>
- Halbheer, D., Stahl, F., Koenigsberg, O., & Lehmann, D. R. (2014). Choosing a digital content strategy: How much should be free? *International Journal of Research in Marketing*, 31(2), 192-206.

- Heiman, A., McWilliams, B., Shen, Z., & Zilberman, D. (2001). Learning and Forgetting: Modeling Optimal Product Sampling Over Time. *Management Science*, 47(4), 532-546.
- Ho, T.-H., Li, S., Park, S.-E., & Shen, Z.-J. M. (2012). Customer Influence Value and Purchase Acceleration in New Product Diffusion. *Marketing Science*, 31(2), 236-256.
- Hoch, S. J. (2002). Product Experience is Seductive. *Journal of Consumer Research*, 29(3), 448-454.
- Jain, D., Mahajan, V., & Muller, E. (1995). An Approach for Determining Optimal Product Sampling for the Diffusion of a New Product. *Journal of Product Innovation Management*, 12(2), 124-135.
- Jing Xu, A., & Wyer Jr, R. S. (2010). Puffery in Advertisements: The Effects of Media Context, Communication Norms, and Consumer Knowledge. *Journal of Consumer Research*, 37(2), 329-343.
- Manchanda, P., Xie, Y., & Youn, N. (2008). The Role of Targeted Communication and Contagion in Product Adoption. *Marketing Science*, 27(6), 961-976.
- Marks, L. J., & Kamins, M. A. (1988). The Use of Product Sampling and Advertising: Effects of Sequences of Exposure and Degree of Advertising Claim Exaggeration on Consumers' Belief Strength, Belief Confidence, and Attitudes. *Journal of Marketing Research*, 25(3), 266-281.
- Mehta, N., Chen, X. J., & Narasimhan, O. (2008). Informing, Transforming, and Persuading: Disentangling the Multiple Effects of Advertising on Brand Choice Decisions. *Marketing Science*, 27(3), 334-355.

- Nam, S., Manchanda, P., & Chintagunta, P. K. (2010). The Effect of Signal Quality and Contiguous Word of Mouth on Customer Acquisition for a Video-on-Demand Service. *Marketing Science*, 29(4), 690-700.
- Narayanan, S., & Manchanda, P. (2009). Heterogeneous Learning and the Targeting of Marketing Communication for New Products. *Marketing Science*, 28(3), 424-441.
- Narayanan, S., Manchanda, P., & Chintagunta, P. K. (2005). Temporal Differences in the Role of Marketing Communication in New Product Categories. *Journal of Marketing Research*, 42(3), 278-290.
- Pauwels, K., & Weiss, A. (2008). Moving from Free to Fee: How Online Firms Market to Change Their Business Model Successfully. *Journal of Marketing*, 72(3), 14-31.
- Sriram, S., Chintagunta, P. K., & Manchanda, P. (2014). Service Quality Variability and Termination Behavior. *Management Science*, *Forthcoming*.
- Su, X. (2009). A Model of Consumer Inertia with Applications to Dynamic Pricing. *Production and Operations Management*, 18(4), 365-380.
- Train, K. E. (2009). *Discrete Choice Methods with Simulation* (2nd ed.). Cambridge, UK: Cambridge University Press.
- Wertenbroch, K. (1998). Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice. *Marketing Science*, 17(4), 317-337.
- Zhang, J., & Niu, B. (2014). Dynamic quality decisions of software-as-a-service providers based on customer perception. *Electronic Commerce Research and Applications*, 13(3), 151-163.

Table 1: Descriptive Statistics

Adoption among free triers and other consumers		
	N	Adopters of the paid service (% of N)
Complete sample	10,000	2,074 (20.74%)
Free triers	1,524	923 (60.56%)
Other consumers	8,476	1,151 (13.58%)

Explanatory variables

	Mean	Total variance	Consumer variance	Time variance
Advertising (€per capita)	.202	.033	.001	.032
Direct marketing (contacts)	.206	.370	.077	.293
Usage during subscription months (00zaps)	1.572	2.604	1.753	.852
Word of mouth (distance-weighted number of surrounding users, 000s)	.342	.134	.035	.099
Onsite repair interventions (% of customers)	3.437	2.081	.085	1.996
Discount (00€)	2.105	.588	.007	.581
Age (years, at start of observation period)	48.516	133.676	133.676	.000
Annual income (000€ average of census block)	25.453	35.180	35.180	.000
Household size (members)	3.111	2.055	2.055	.000
Relationship length (years since DSL adoption)	2.215	1.516	1.516	.000
Advertising main competitor (€per capita)	.072	.008	.004	.003
Trend (months since launch)	11.96	46.430	2.079	44.351

Table 2: Comparison of Alternative Model Specifications

Model	Nr. of parameters	Log Likelihood	BIC	AIC	Root mean square error ^a	
					In- sample	Out-of- sample
BM1	24	-15,368.03	31,030.93	30,784.06	5.385	5.181
BM2	33	-14,726.85	29,859.16	29,519.70	4.286	4.180
BM3	35	-14,748.42	29,926.87	29,566.84	4.197	4.108
BM4	37	-14,364.06	29,182.73	28,802.13	4.461	4.174
FM	40	-14,354.95	29,201.35	28,789.89	4.082	4.035

^a in thousandths

Table 3: Parameter Estimates

	Estimate (standard error)	
	Population Mean	Standard Dev.
Coefficients of utility functions		
<i>Core coefficients</i>		
Intercept, β_i^0		
Constant, α^0	6.522 (.788) ^a	.001 (.137)
Age (years, at start of observation period)	-.015 (.002) ^a	
Average annual income (000€)	-.009 (.005) ^c	
Household size (members)	.056 (.019) ^a	
Relationship length (years since DSL adoption)	-.009 (.022)	
Free-trial constant, β_i^F	7.231 (.353) ^a	.038 (.176)
Discount (€), β_i^{DISC}	.004 (.001) ^a	.001 (.001) ^b
Risk aversion w.r.t. quality beliefs, $r_i = \exp(r_i^*)$ (lognormal)	3.192 (.956) ^{a, d}	.574 (.187) ^{a, d}
Uncertainty reduction, β_i^{UR}	28.821 (16.101) ^c	.215 (.796)
Fixed effect after free trial, β_i^{FE}	4.692 (.166) ^a	.031 (.235)
<i>Coefficients of remaining control variables</i>		
Own advertising (€per capita)	.308 (.163) ^c	.077 (.398)
Direct marketing (contacts)	.275 (.032) ^a	.398 (.062) ^a
Advertising main competitor (€per capita)	.194 (.277)	.007 (.637)
Trend (months since launch)	-.460 (.031) ^a	.004 (.005)
Paid-trial constant	4.089 (.449) ^a	1.116 (.395) ^a
Fixed effect after paid trial	6.202 (.481) ^a	.459 (1.373)
Parameters of learning process		
Consumers' initial expected ultimate match quality, \bar{q}_0	-.594 (.176) ^a	
Ultimate mean match quality, θ_i	5.195 (1.649) ^a	.005 (.017)
<i>Match quality signals' standard deviations</i>		
Std. dev. of advertising signal, σ^A	.034 (.003) ^a	
Std. dev. of direct marketing signal, σ^D	.023 (.001) ^a	
Std. dev. of usage signal, σ^U	.203 (.007) ^a	
Std. dev. of WOM signal, σ^W	.037 (.004) ^a	
<i>Signal bias and bias beliefs</i>		
Mean overcorrection of marketing signals, $\bar{d}^* = \delta - \bar{d}$	-5.367 (1.690) ^a	
Std. dev. of belief about marketing signal bias, s^d	.000 (.052)	
Real mean match quality fluctuation coefficient, κ	-2.078 (.649) ^a	
Mean of belief about real match quality fluctuation, \bar{k}	-.014 (.007) ^b	
Std. dev. of belief about real match quality fluctuation, s^k	.023 (.006) ^a	
Log likelihood value: -14,354.947; BIC: 29,201.354; AIC: 28,789.894		
$2 \times (\text{Loglik full model} - \text{Loglik homogenous intercepts only}): 11,322.066 (p < .001)$		

^a Significant at the $p < .01$ level; ^b Significant at the $p < .05$ level; ^c Significant at the $p < .10$ level;

^d Standard errors are derived from the covariance matrix of the estimated parameters μ_{r^*} and σ_{r^*} , using the delta method.

Figure 1: Consumers' Decision Process in the Presence of a Free-Trial Promotion




Figure 2: Impact of a Free-Trial Promotion on Adoption: A Decomposition

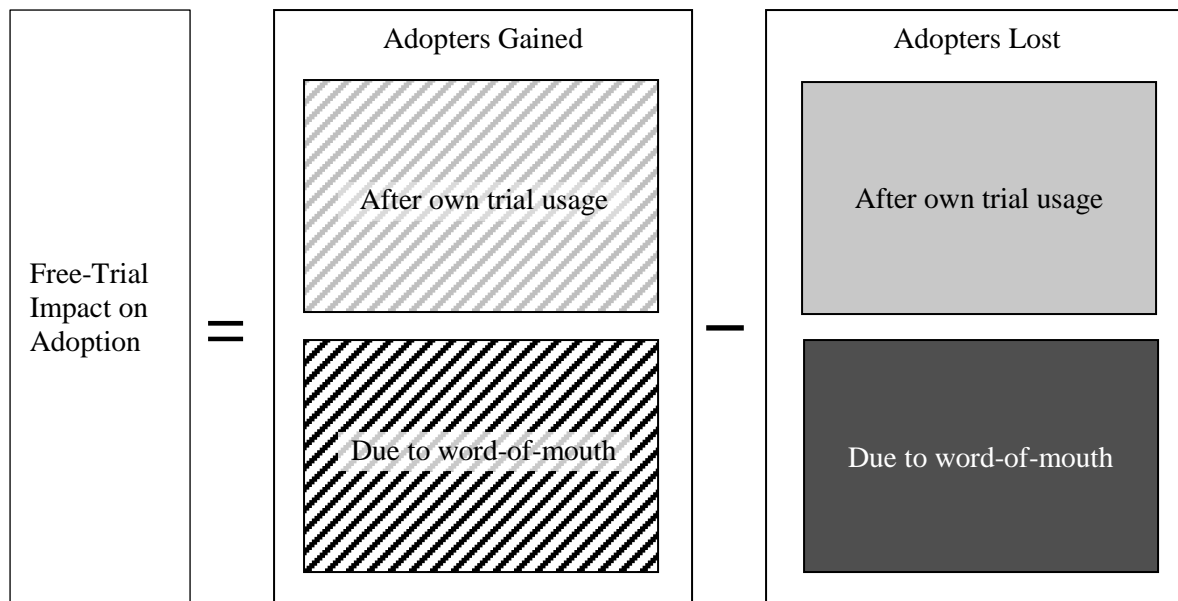


Figure 3: Illustration of a Consumer's Varying Choice Set




Figure 4: Observed and Predicted Free-Trial Subscription and Subsequent Adoption over Time




Figure 5: Uncertainty About Ultimate Match Quality Over Time




Figure 6: Average Corrected Quality Signals Over Time

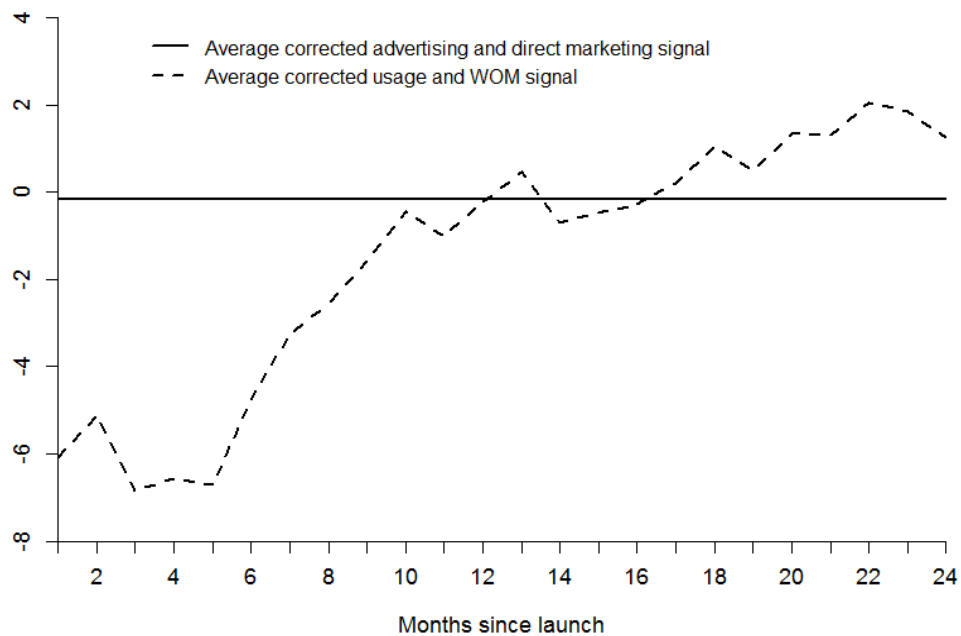


Figure 7: Simulated Free-Trial Effect: Number of New Adopters Over Time




Figure 8: Simulated Impact of Timing on Total Number of Adopters




Figure 9: Simulated Impact of Timing on Total Number of Paid Adoption Months

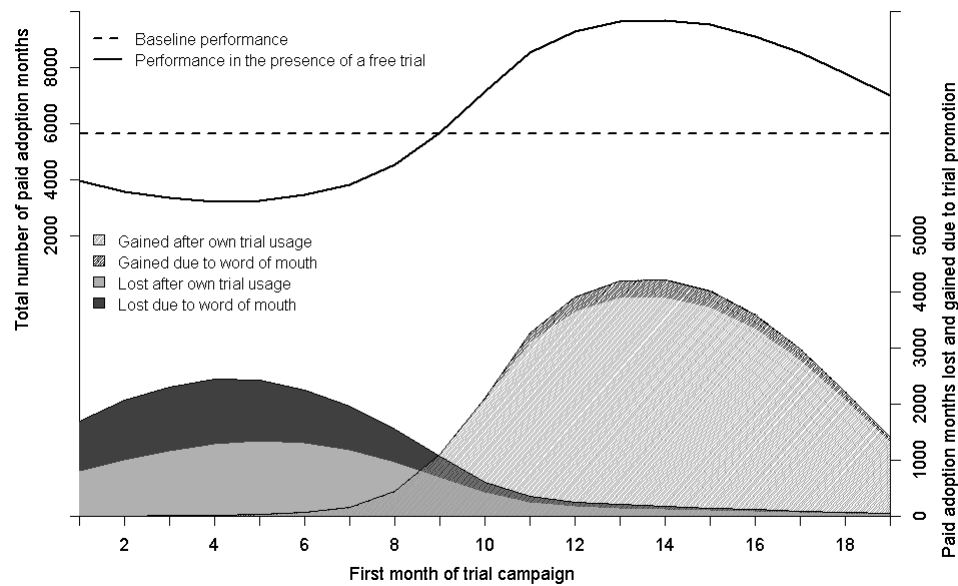



Figure 10: Simulated Impact of Usage Intensity on Total Number of Adopters

a) Trial Campaign Starting in Month 17



b) Trial Campaign Starting in Month 5

