# New Features Free of Charge?
# Using Price to Sort Consumers Among Legacy Software Versions

Kristina Brecko*

September 16, 2016

PRELIMINARY AND IN PROGRESS
Current Version: http://stanford.edu/∼brecko

**Abstract**

In many durable good contexts, firms have the opportunity to price discriminate on quality by charging higher prices for the latest functionality. In the software good market, on the other hand, we often do not observe price discrimination on the latest versions, despite new versions being introduced over time. I propose that the software firm's ability to price discriminate on latest functionality is restricted by two factors: (1) the extent to which consumers value the innovation from one version to the next and (2) the extent to which legacy software products are costly for the firm to maintain. To analyze this question, I use a unique dataset on individual consumer subscriptions to a Fortune 500 firm's software products. The firm releases new product versions each year, but allows consumers to adopt the latest functionality for free. Despite this policy, descriptive analysis reveals that consumers frequently choose not to upgrade, electing to renew legacy versions of the product instead. To distinguish between the different factors driving this pattern, I develop a dynamic model of consumer choice of different product versions, renewal opportunities and upgrades. This model allows me to separately account for version usage utility, non-monetary costs of purchasing and upgrading and the heterogeneity therein. The estimates of the model reveal that although the majority of the consumers value the new versions, the high value, price insensitive consumers do not, causing it to be unprofitable for the firm to price latest functionality at a premium. Using the estimates and the structure of the model, I further describe a counterfactual that allows me to quantify how much a firm must innovate in order to be able to price new functionality at a premium when legacy versions are costly. The final counterfactual allows me to calculate the minimum legacy version cost that would cause the firm to shift from releasing distinct intertemporal versions to maintaining one continuously upgraded version of the product.

*Doctoral Student in Marketing, Stanford Graduate School of Business, Email: brecko@stanford.edu. I would like to thank my advisor Wesley Hartmann for his guidance and support. A special thank you also to Peter Reiss for his invaluable advice. I am also grateful to Tomomichi Amano, Timothy Bresnahan, Pedro Gardete, James Lattin, Harikesh Nair, Sridhar Narayanan, Navdeep Sahni, Stephan Seiler, Anna Tuchman, Ali Yurukoglu and the participants of the Stanford IO and Marketing seminars for their comments. I thank the Wharton Customer Analytics Initiative and an anonymous sponsoring firm for the data resources. All remaining errors are my own.

# 1 Introduction

Firms often price discriminate on quality by introducing different versions of the same product. In the typical second degree price discrimination example, a firm may introduce an expensive product version with the best functionality alongside a cheaper product version with lower-level functionality, such that each consumer sorts into a particular quality tier based on his preference for quality and willingness to pay (Varian 1997, and more recently, Bhargava and Choudhary 2008). Firms selling durable goods with evolving product quality (e.g., cameras) have the opportunity to use innovation over time as an additional price discrimination dimension. By setting a high price for a product version with the latest functionality, the firm allows consumers to sort accordingly: high valuation, high willingness-to-pay consumers purchase the latest version and everyone else continues to use older versions of the product. Software products are another example of a durable good with evolving product quality, and another context in which we may expect to observe this type of price discrimination; however, in software markets, we often do not observe higher prices for the latest versions, despite new versions being introduced over time. In fact, in many cases, software firms offer new functionality at a discount, thus, giving consumers an additional incentive to upgrade to the newest version. Moreover, via changes in pricing and product design, software firms are increasingly shifting away from maintaining a number of intertemporal product versions, thus eliminating altogether the ability to price discriminate on new functionality.

One extreme example of a firm making this shift is Microsoft's 2015 decision to offer Windows 10 as a free upgrade for Windows 7 and Windows 8 users until July 2016 (Microsoft 2015). By doing so, industry specialists argue that Microsoft forgoes any possible gains to new version sales in favor of the gains from managing one version of the operating system[1] (Hoffman 2015). Other software firms have made similar pricing and product design changes. For instance, in 2013, Adobe discontinued perpetual licenses to intertemporal product versions of the Adobe Creative Suite, in order to accelerate a shift over to the subscription-based Adobe Creative Cloud, a one-version software product that receives free latest-functionality upgrades (Adobe 2013). Similarly in 2011, Microsoft introduced Office 365 as a subscription-based product with free updates, alongside intertemporal versions of Microsoft Office products (Microsoft 2011). These examples are all part of a larger cloud-computing trend in the software industry away from distinct intertemporal product versions and towards a subscription-based model with one product version that receives free upgrades to the newest functionality over time (Knowledge@Wharton 2013). These changes leave us with a puzzle: for many durable goods with evolving product quality we see firms using new functionality as a price discrimination tool, and yet in software settings we do not. Moreover, firms are increasingly eliminating their ability to price new versions differently by no longer releasing distinct product versions over time.

In this paper, I identify two factors that restrict the software firm's ability to price discriminate on the latest functionality and that explain why a firm might eliminate legacy versions altogether.

[1] Such costs may include costs of i) tailoring application development and security patches towards disparate set of legacy versions and ii) providing timely operating system support.

The first factor is consumer valuation of successive product versions. The proportion of the firm's most valuable consumers and the extent to which each values each successive version affect the firm's incentives to charge a higher price for the newest functionality. For example, the firm may not have an incentive to price the new version at a premium if the high value consumers value old functionality more than new, while low value consumers prefer the functionality of the new version. This particular demand-side explanation may be especially relevant for mature products, such as word processors, where incremental innovation provides improved look or speed of the product rather than drastically different functionality. It is also particularly important in operating system and security software contexts, where firms often further manage the use value of the software via intra-version security patches. In this case, the changes from one major product version to the next may not be disruptive enough for the existing high value consumers to value the new version more than an older version they know and like. Thus, if the high value consumers comprise a sufficient portion of the population, the firm does not see an incentive to price discriminate on successive versions.

The second factor concerns the firm's costs associated with maintaining legacy versions of the product. In the typical durable good example, the firm has limited costs associated with a product after it is sold. The same is not the case with software products, where there is often an expectation that, even after the sale, the software firm will continue to maintain the product (e.g., via security patches) and make it compatible with new operating systems and other new software products. This particular cost-side factor is unique to software, but varies with the extent to which the firm is expected to provide this type of support for its in-use products. For instance, the cost may be very high for an operating system like Microsoft Windows, but lower for a content creation suite like Microsoft Office. From the firm's perspective, legacy software versions can be more difficult to service than current versions due to specialized manpower needed to deal with legacy code, for instance. While software firms may not incur these costs if only a few consumers use legacy versions, they incur greater costs with higher probability the more consumers use legacy versions of the product. Moreover, legacy versions may be more vulnerable to piracy, further contributing to costs. As a result, to discourage legacy version use, a software firm with high legacy version costs may be reluctant to set higher prices for the latest functionality. In fact, if these costs are high enough (Informatica 2013), it may choose to retire legacy versions altogether or shift away from releasing distinct product versions over time.

Using novel data from a Fortune 500 software firm, I empirically test whether the nature of consumer valuation of successive product versions can explain the observed price discrimination on new functionality. I find that given the heterogeneity in consumer valuations of each successive version, it is, indeed, optimal for the firm to price the older version at a premium, and offer upgrades for free. I then consider other product design scenarios where it is profitable for the firm to price new versions at a premium. In this context, I design a counterfactual analysis to show that even in cases when pricing new versions at a premium is profitable based on consumers' valuations, increased costs of legacy versions may prevent the firm from setting higher prices for the newest version. Finally, I describe a counterfactual analysis that allows me to calculate how costly legacy versions must be in order for the firm to shift away from releasing distinct versions over time, towards one version that is continuously maintained and upgraded. The focus of this paper is to examine the limits of traditional price discrimination on new functionality in a context with limited incremental innovation and costly legacy versions. The results of this paper further help characterize and explain the trend in software away from distinct intertemporal versions.

In my application, I observe lower, rather than higher, prices for newer versions of the product and free upgrades to the newest functionality. Yet, I observe many consumers still using legacy products. First, I present descriptive evidence to establish the main factors driving this consumer behavior. Important factors in this context include monetary and non-monetary costs of purchasing, renewing or upgrading to the latest version of the product, as well as heterogeneity in consumer valuations of the versions. While these factors are intertwined, constructing and estimating a model of consumer demand allows me to separately identify these different elements. To measure the primitives of consumer behavior driving product demand, I develop and estimate a dynamic model of consumer choice. Future opportunities to purchase, renew or upgrade the product to the newest version inherently introduce a dynamic component into the consumer's current-period choices. Following the literature on durable good purchase and use decisions, I allow for forward-looking consumer behavior with respect to subscription length, version and version release timing.

I find that consumers who most value the product overall value the new version less than the previous one, while low-value consumers seem to value incremental innovation. This result is opposite to the one documented in previous work on the impact of innovation on video game usage by Albuquerque and Nevskaya (2015). Moreover, the result is the flip-side of the innovator's dilemma described by Christensen (1997), where the firm loses out on a new sector of demand by tailoring the product to its most valuable consumers. In this case, rather than suffering from this innovator's problem, the firm seems to be tailoring the product for broader consumption to the detriment of its most price insensitive consumers. Using the estimated demand parameters, I further evaluate whether the firm is pricing optimally. I construct an illustrative model of the firm's pricing decisions across versions. The results of this model are roughly consistent with the observed pricing scheme - that under the given demand parameters, it is optimal for the firm to set a higher price for the older versions and offer the newer versions at a discount.

Through a counterfactual exercise, I evaluate a case where high value consumers value the innovation more than low value consumers. The illustrative model suggests that in this case the firm should price new versions at a premium and charge the consumers to upgrade to the newest version. I show that, indeed, the firm's profits in this case are higher under this price discrimination scheme than they would be under the current pricing schedule. However, I describe a counterfactual in which I can show that these profits are eliminated with costly legacy versions. Finally, I consider the case where the firm no longer allows any legacy versions, but instead moves all consumers on to the most recent version as it becomes available, an outcome that mimics the shift we are seeing in the software industry. In this case, the firm loses some profits from dissatisfied high value consumers, but gains by eliminating legacy versions altogether. This counterfactual allows me to characterize how costly the legacy version maintenance must be in order for the firm to make this move.

To my knowledge, this paper provides the first empirical analysis of the factors affecting the shift away from intertemporal price discrimination via versioning. Most of the analysis is focused on understanding the heterogeneous demand for versions of the software product and how that affects the firm's ability to price discriminate; however, by analyzing and estimating the demand model, I can also evaluate cost-side constraints that further inform the firm's pricing policy of new versions. Thus, this paper contributes to our understanding of intertemporal product version and upgrade pricing, in contexts with little incremental innovation one version to the next and costly legacy versions.

In the sections that follow, I position this paper in the past literatures on product version price discrimination, durable good and innovation pricing, and release timing. I then describe my particular empirical context and data set. After presenting a descriptive analysis of the data, I construct the dynamic model of consumer choice and the variation in the data that identifies the model parameters. After a brief discussion of the estimation approach, I present the results. I then use the estimation results to evaluate the impact of counterfactual product designs and legacy costs on pricing decisions. I conclude the paper by summarizing my findings and detailing directions for future research.

## 2 Literature Review

This paper relates to literatures on product version price discrimination, innovation pricing and release and, more broadly, durable good pricing. In this section, I briefly discuss these literatures, and how my paper contributes to each.

Previous work has analyzed the conditions under which version price discrimination at a given point in time is optimal for an information goods producer (negligible marginal costs of production). Varian (1997) shows how differential quality of versions enables the firm to price discriminate and analyzes welfare implications. Since Varian (1997), a number of marketing scholars have also focused on static product line design and pricing in contexts other than information goods (e.g., analytical studies by Desai, Kekre et. al. (2001), Villas-Boas (2004), Desai (2006), Netessine and Taylor (2007) and empirical studies by Draganska and Jain (2006) and Li (2014))[2]. More recently, Bhargava and Choudhary (2008) derive that, at a given point in time, it is optimal for the information goods firm to introduce a low quality-price product version if the high type consumers have a relatively low valuation for the quality-price version (and vice versa) and the marginal costs of production are sufficiently low.

In this paper, rather than considering the product assortment and version pricing at a given

[2] Analytical work studying product line design at a given point in time include Desai, Kekre et. al. (2001) and Netessine and Taylor (2007), who examine the relationship between manufacturing costs and product line design, Villas-Boas (2004), who explores the relationship between product line design and advertising, and Desai (2001) who considers competitive forces. Empirical studies have also examined the provision and pricing of variety in categories other than information goods. For instance, Draganska and Jain (2006) find that yogurt manufacturers are acting optimally when price discriminating on quality tiers, but not on flavor attributes. Li (2014) considers optimal provision of variety, given recent consumer preference shocks.

time, I consider the pricing of intertemporal product versions for information goods with evolving product quality, which relates more closely to the literature on upgrade pricing, initiated by Dhebar (1994) and Fudenberg and Tirole (1998). Among other analytical results, Fudenberg and Tirole (1998) show that in the information goods context, a firm's upgrade pricing policy depends on the incremental quality improvement from one version to the next (but allow only positive incremental valuation) and the incremental cost of producing the new version. Viard (2007) further builds on this work and demonstrates that the full version price and version upgrade prices are increasing in the information good innovation rate. Moreover, Zhang and Seidmann (2010) propose a model that relates the uncertainty of firm innovation and the software product's network effect to the firm's optimal choice of software product licensing (perpetual license, subscription license or both).

To my knowledge, this paper is the first to conduct an empirical study of the software firm's upgrade release and pricing policy, as a function of consumer heterogeneity and legacy version maintenance costs. Unlike Goettler and Gordon (2011) who analyze the effect of competition on durable goods innovation and its pricing and release, I focus on the firm's ability to price discriminate based on consumer heterogeneity and its own costs[3,4]. Unlike prior empirical literature on durable good replacement and usage (Gordon 2009, Albuquerque and Nevskaya 2015), in addition to studying uptake and usage of new versions, I examine the firm's ability to price discriminate on new functionality as well as conditions when such price discrimination may not be optimal. In addition, I consider how a consumer's upgrade probability is affected by the non-monetary costs of upgrading, such as installing a new product version and learning its new functionality, constructs which are more related to the literature on learning and switching costs (e.g., Narayanan et. al. 2007, Goettler and Clay 2011). I also analyze a more recent trend in software markets away from distinct intertemporal versions, which has not been considered in the literature. I propose that, in addition to the shape of consumer heterogeneity, this shift may be also prompted by higher costs of maintaining legacy versions, a cost structure similar to the one documented in damaged goods literature (e.g., Deneckere and McAffee 1996) and the analytical marketing literature referenced above (e.g., Desai, Kekre et. al. (2001) and Netessine and Taylor (2007)).

This paper relates more broadly to work on durable technology good pricing[5], albeit in a context where a firm offers technology good rental plans. It also relates to the obsolescence literature[6], given that legacy versions of some software products (e.g., security software) effectively become obsolete if the firm discontinues issuing securing patches. Finally, the paper is also related to the extensive marketing literature on innovation and innovation diffusion[7].

[3]An extension of this paper would further endogenize the version quality by modeling the firm's innovation process and incorporating consumer uncertainty about future version quality.

[4]Related theory papers that also study the optimality of version releases are Ellison and Fudenberg (2000) who explore the social optimality of frequent software product version releases and Borkovsky (2015) who builds on Goettler and Gordon (2011) by making the firm's release timing of version releases endogenous.

[5]Among others, Nair (2007) in the video game context and Gowrisankaran, Rysman (2012) in the camcorder context.

[6]Starting with Bulow (1986), Levinthal and Purohit (1989), then Fishman and Rob (2000).

[7]See Hauser et. al. (2006) and Peres et. al. (2010) for a review.

# 3 Empirical Context and Data Discussion

## 3.1 Data Description

The data used in this study come from a Wharton Customer Analytics Initiative (WCAI) data grant, provided courtesy of an anonymous software firm. The firm provided four different data sets: consumer data, purchase and renewal data, price data, and "heartbeat" data, used to infer upgrading behavior. The consumer data tracks nine cohorts of the firm's consumers, where a consumer is identified by an account he creates with the firm, and a consumer's cohort is defined by the year of his first purchase of one of the firm's products (2006-2014). The cohorts are equally-sized random samples of all of the firm's personal and business account consumers[8] first purchasing in a given year, such that each consumer cohort comprises about 11% of the total consumer population.

The second data set records consumers' purchase and renewal activity. The firm sells mostly 12 and 24 month subscriptions to three major software products (Product A, Product B, Product C) and a number of smaller products, grouped into an "Other" product group. The firm tracks the consumers' subscriptions at the license level. For instance, the consumer can purchase a 12-month subscription license for Product A and, when the subscription expires, renew this existing license or purchase a new license (through a different channel, for example). In my descriptive analysis and estimation, I model the new purchase and renewal choices separately.

The third data component contains the price paid upon purchase of a new subscription license or renewal of an existing product license. The price is paid in 12 or 24 month intervals, only upon purchase or renewal, and differentiated by four factors: product, purchase vs. renewal, subscription length and month. The firm and product context is anonymous, and the price data is normalized; however, in what follows, I use the dollar unit ($) when referring to the prices for ease of exposition. Free software trials of different lengths are also available and are more prevalent for some products (Product C and, to a lesser extent, the Other product group) than others (Products A and B)[9].

The fourth data component contains "heartbeat" data on each subscription license. Whenever there is a change in the configuration of the software license, the firm receives a "heartbeat" from the system on which the software is installed. The "heartbeats" tracked in this data set include upgrades to the new product version, changes in auto-renewal enrollment and installation on a different machine, among others. I use the "heartbeat" data to supplement the renewal and purchase data with consumer upgrading behavior from one major product version to the next. With a few exceptions, the firm introduces new major product versions at an annual rate[10]. All consumers who have a valid subscription are eligible to upgrade to the newest version of the product for free. In order to complete an upgrade, the consumer has to either click on a button on the software interface or go directly to the firm's website.

[8]Excludes enterprise software product licenses.

[9]See Table 1.

[10]In addition to the major product versions, the firm also introduces intra-version security updates, which consumers receive automatically.

## 3.2 Data Preparation and Population Selection

To arrive at a data set suited for descriptive analysis and estimation, I undertake several steps to merge the four data components described above and apply several population selection criteria. In this section, I give an overview of the key decisions made to arrive at the final data set[11].

**Study Population** First, I restrict my attention to personal account consumers. These constitute $\approx$ 84% of the total population of consumers without any missing or inconsistent purchase and renewal data[12]. I exclude small business licenses because they are typically geared towards use by multiple people in a work setting. Therefore, the utility function of the small business users is likely to be very different from the utility function of the individual personal consumers. Secondly, I focus only on the 2009-2013 consumer cohorts, representing $\approx$ 72% of the personal account consumers described above, for several reasons. The earlier cohorts (2006-2009) disproportionately have missing or inconsistent data and arrive during a period when neither price paid nor product version / upgrade data are available (price data are available starting in January 2010 and version information is available starting in January 2009). Since price and version information are crucial to my analysis, dropping the early cohorts and their activity does not limit my data panel much, but ensures that the remaining cohorts are represented approximately equally. I drop the 2014 cohort because the data set contains, at most, half a year of activity for these consumers. Finally, from this set of 2009-2013 consumers, I keep only those who purchase 12 and 24 month subscription licenses in the January 2009 to June 2014 time-frame ($\approx$ 87% of cohort consumers), are not seen to downgrade products or upgrade expired products ($\approx$ 98% of cohort consumers) and are not seen to use multiple machines simultaneously ($\approx$ 88% of cohort consumers). I abstract away from consumers who appear to be managing subscriptions on a number of different machines in order to keep the analysis at a consumer level rather than a consumer-machine level.

**Master Data Set of Consumers' Purchases, Renewals and Upgrades** As a first step, I merge the consumer data with purchase / renewal data and the price data to obtain a consumer-level data set of purchases and renewals with the corresponding prices paid. As a second step, I add information about the version being used at a given point in time. As discussed above, the upgrade information is contained in the "heartbeat" data. That is, the consumer's system sends a "heartbeat" to the firm whenever the consumer upgrades to the newest version. Since purchases, renewals and upgrades often occur on different months, it is necessary to infer the version information from the upgrading activity in order to understand which version of a given product is purchased or renewed. Inferring the version is straightforward after the first upgrade or other "heartbeat" activity. Prior to the first upgrade, I infer the product version in the following manner. If the new subscription license purchase date is close to the first "heartbeat" date (within 3 months) and the first "heartbeat" registers the use of a legacy product, I infer that the consumer purchased the a subscription license

[11]The comprehensive data preparation and population selection description is available upon request, conditional on sharing approval by the sponsoring firm.

[12]Consumers with missing or inconsistent data include those for whom the recorded subscription start date is after the expiration date, for whom the renewal date is missing or for whom the the first purchase is not observed.

for the legacy product version shown in the "heartbeat" observation. For the remainder of the observations, I assume that the consumer purchased a subscription to the newest available version of the product. This assumption is not restrictive, since in most cases where product version information is available on the purchase date, new license purchases are for the newest available version. Moreover, many of the subscription licenses are used in their initial version. Thus, any error in the version inference will not affect my upgrade analysis[13].

**Price Panel** As is common in analysis of consumer purchase data, I only observe prices that were paid upon purchase or renewal of a subscription to a particular product. I construct a panel of prices that a consumer faced when making a purchase decision in two steps. First, as discussed in the previous paragraph, I obtain a customer-level data set of purchase and renewals and the corresponding prices paid. In order to do that, I merge the consumer purchase and renewal data with the price dataset; however, the match is not one-to-one; i.e., the price paid is missing for some renewals and purchases. Next, I fill in the missing prices for every transaction. I first try to match an observation with a missing price to an observed price paid on the same month for the same sku[14], subscription length, purchase or renewal category combination[15] (sku-level match). In the absence of a sku-level match, I try to match an observation with a missing price to an observed price paid on the same month for the same *product*, subscription length, purchase or renewal category combination (product-level match). Finally, in absence of a product-level match in the same month, I infer the price paid from the nearest (in time) sku or product-level match ($\sim$ 22% of the price observations interpolated in this manner). In the second step, I construct a time panel of median prices by product, subscription length, and purchase or renewal category by (1) taking the median of the prices observed and inferred in Step 1 for the given attribute combination, and (2) interpolating the price in the manner described above for months with no observed prices paid for the given attribute combination. Figures 9 and 10 in Appendix Section A graph the price panels for the three main products offered by the firm. The two-step approach works well even if the price data are not missing at random, since it does not underweight the purchases with missing prices; however, it might introduce measurement error due to the interpolation. In future work, I plan to check robustness of the obtained results by re-running the estimation procedure with median prices obtained from a one-step approach (using the product-level panel which takes into account only observed prices) and 25th and 75th percentile prices rather than median prices.

As a result of this preparation process, I arrive at a consumer-month level data set of $38,086$ consumers from the 2009-2013 cohorts. These consumers have 84,991 observed purchases and renewals (trial and paid) from January 2009 to June 2014. Most overall purchase and renewals are for Product C, however, a large portion of these are free trials. On the other hand, Product A has the highest number of paid purchases and renewals and, correspondingly, the highest number of

[13]By inferring the version information in this manner, I may overstate the number of newest version products in the market relative to legacy products; however, this bias works against any results related to the total cost of legacy version usage to the firm.

[14]Sku stands for stock keeping unit, which is a more granular categorization than "product."

[15]If there are several matching price observations, I take the median price.

Table 1: Study Population - Consumers, Licenses Purchases & Renewals and Upgrades by Product

| | **Consumers** | **Licenses** | **Purch & Renew** | **Upgrades** | **Purch & Renew** | | **Paid Purch & Renew** | |
|---|---|---|---|---|---|---|---|---|
| | | | | | Trial | Paid | 12 Mo | 24 Mo |
| Product A | 14,081 | 16,364 | 26,877 | 9,540 | 1,037 | 25,840 | 24,764 | 1,076 |
| Product B | 7,538 | 8,390 | 11,498 | 2,173 | 528 | 10,970 | 10,777 | 193 |
| Product C | 19,436 | 24,141 | 40,606 | 8,333 | 14,873 | 25,733 | 24,791 | 942 |
| Other | 5,495 | 5,779 | 6,010 | 0 | 4,310 | 1,700 | 1,700 | 0 |
| Total | 38,086 | 54,674 | 84,991 | 20,046 | 20,748 | 64,243 | 62,032 | 2,211 |

Note: "Purch & Renew" refers to all new subscription license purchases and existing license renewals (paid and trial).

product upgrades. Twelve month subscriptions are predominant, with 97% of all paid transactions being for twelve month subscription licenses.

# 4 Descriptive Evidence

In this section, I present descriptive analysis of consumer purchase, renewal and upgrade behavior, elements that motivate the model in Section 5. Throughout most of this section, I focus on Product A[16]. The patterns for products B and C are similar (with minor differences mentioned below) and can be found in Section B of the Appendix.

## 4.1 Purchase and Renewal Behavior Patterns

**Aggregate Purchase and Renewal** The first pattern I document is that consumers buy the latest version of the product when purchasing a new subscription, but continue to renew existing subscriptions for legacy product versions. This pattern is reflected in Figure 1. On the top panel of the figure, I plot the aggregate number of new Product A subscription license purchases by version by month, highlighting, in particular, version 5 of the product. The top panel shows that during the time period between the release of two versions, the vast majority of new subscription purchases are for the newest version of the product (e.g., version 5 during the February 2011 - January 2012 time period)[17].

Existing subscription license renewal activity, however, follows a different pattern. The bottom panel of Figure 1 shows that even when a newer version of a product is released, those consumers who renew existing subscriptions continue to do so without upgrading to the newest product version. For example, even after version 6 is released in February 2012, consumers continue to renew version 5 of the product. In fact, a non-negligible portion of the consumers continue renewing the subscription for product version 5 even towards the end of the panel in 2014. Thus, in this applica-

[16]For reasons discussed in Section 6, I use Product A purchase, renewal and upgrade data for estimation, and, as a result, the descriptive patterns for this product are especially relevant.

[17]This pattern is somewhat weaker for Product C, as shown in Figure 11 in Appendix Section B. The reason for this may be that a trial version of Product C frequently comes pre-installed with the purchase of a machine (observed in the data). The pre-installation may create a lag, and a new version of the product may become a legacy version of the product by the time a machine makes its way from assembly to a consumer's home.

## Figure 1: Paid New Purchases and Renewals by Version (Product A)

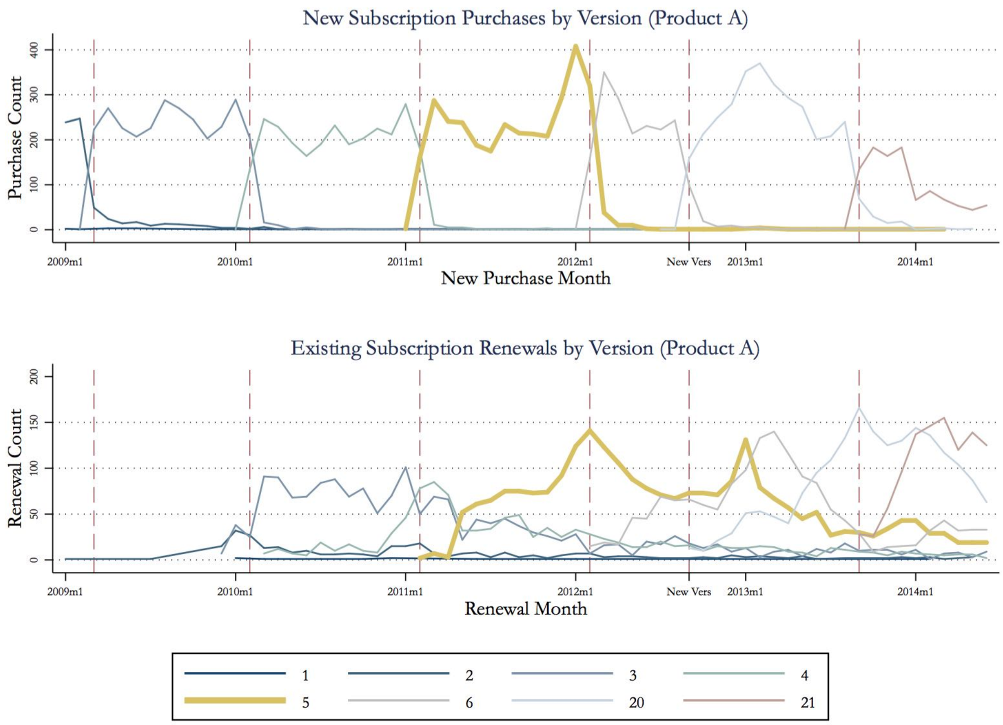


Note: graph includes paid purchases and renewals only, excludes trials.

## Figure 2: Upgrades to a Given Product Version (Product A)

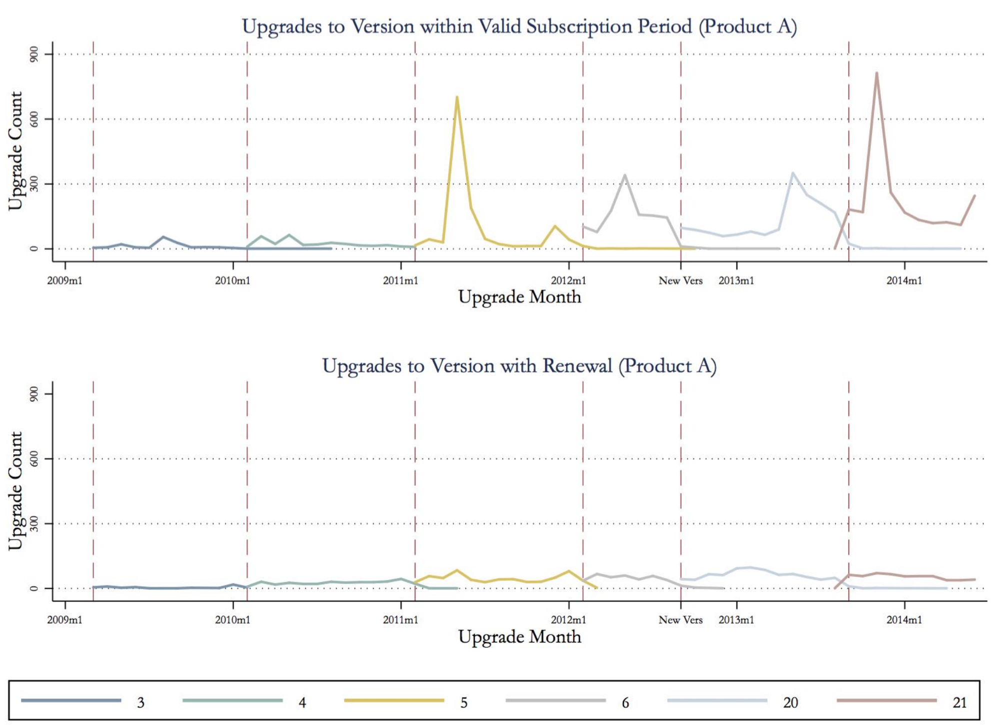

tion, consumers frequently choose not to upgrade their software products to the newest version, and legacy version usage results from renewals of existing subscription licenses without upgrades rather than purchases of new licenses of legacy versions. There are several possible drivers of this pattern, including state dependence, differences in version valuation, or heterogeneity in these elements. In the following paragraphs, I examine the data further for evidence of these different forces.

**Variation in Reactivation by Version** To better understand heterogeneity in re-purchase and renewal behavior, I further examine the relationship between version reactivation (re-purchase and renewal) decisions and reactivation prices paid (as proxy for valuation). Focusing on consumers who first purchased a particular version of the product, version 4, and their first reactivations of the subscription license, in Figure 3 I plot the distribution of product versions at the time of first reactivation for the cohort of consumers who start on the same version. The left panel shows that of those consumers who first purchased version 4, consumers who have a lower valuation for the product, as indicated by the price paid, are more likely to reactivate the product in the newer versions 5 and 6. On the other hand, consumers who have a higher valuation for the product are more likely to reactivate the product in the version on which they started.

In Appendix Section H, I present a simple model of the firm's pricing decisions when two types of consumers (high and low) have different valuations for a new and old version of the product. In this simple case, in order for the firm to profitably charge higher prices for newer versions, the difference in the high type's valuation of the two versions must be sufficiently large and higher than the low type's (i.e., single crossing condition should hold). In this case, the firm can use price to sort the consumers: with high type consumers purchasing the new version at a higher price and the low type consumers purchasing the old version at a discount. The descriptive pattern in Figure 3 reveals the opposite trend - one where the high value consumer at a higher rate renews older versions of the product - suggesting that a necessary condition for new functionality price discrimination may not be satisfied.

The figure does not reveal, however, whether the pattern results from the high type consumer's lower valuation for new versions or because of state dependence considerations and heterogeneity therein. In the following sub-section, I further discuss the data patterns that speak to both of these elements.

## 4.2 Upgrade Behavior Patterns

In this application, a consumer can upgrade whenever he has a valid product subscription and is using a legacy version of the product. Since the reactivation decision is de-coupled from the upgrade decision, a consumer's choice of whether or not to upgrade when eligible further informs my understanding of his version valuation and state dependence.

**Variation in Version Uptake Rates** Table 2 shows the rate at which consumers who are eligible for an upgrade, i.e., have a valid subscription license and are using a legacy version of the

Figure 3: First Time Re-purchases and Renewals by Price Paid (Product A)

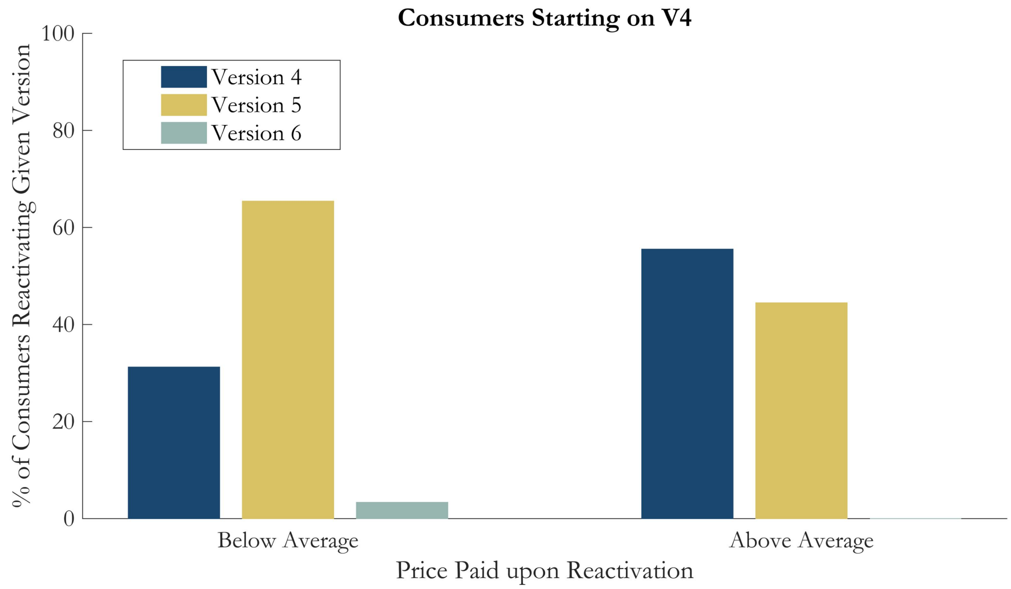


Note: This descriptive pattern employs the data used for estimation. For a discussion of the selection criteria applied, see Section 6.

product, upgrade from one version to the next. The rows represent the legacy version of the product from which the consumer is upgrading, and the columns represent the snapshot in time when a particular version is the latest available product version. The highlighted fields represent the rate at which eligible consumers upgrade to version 20 of the product from the different legacy versions. Several patterns emerge.

First, the rate at which eligible consumers upgrade to the newest version is decreasing with the distance from that newest version. For instance, upgrade-eligible consumers using legacy version 3 of the product are much less likely to upgrade to version 20 than upgrade-eligible consumers using legacy version 6. One reason for this trend could be that upgrading across multiple versions is more complicated and costly than upgrading to the next version. The functionality changes from version 3 to version 20 are a product of four years worth of innovation, and a consumer upgrading from version 3 to 20 has to learn to use the new functionality features all at once. Moreover, because of multiple functionality changes, an upgrade across several versions may introduce more software-machine incompatibilities that would need to be resolved after the upgrade than an upgrade to the next version. Another explanation for this behavior, however, is that there is something inherently different about consumers who have continued using the product in version 3 four years after this version was first released. It could be that these consumers have a high use value for version 3. Alternatively, it could be that they have particularly high costs of upgrading because they highly value the uninterrupted use of their computer. Although I am unable to distinguish these two explanations here, I account for these two forces in the structural model in Section 5.

A second notable pattern from Table 2 is that the overall rates at which eligible consumers upgrade to a particular newest version vary by product version. While 28% of eligible consumers upgrade to version 6, 38% of eligible consumers upgrade to version 20 and 43% of eligible consumers upgrade to version 21. The low rate of upgrade to version 6 is partially a result of the shift in Product

Table 2: Proportion of Eligible Consumers Upgrading (Product A)

| | Newest Available Version | | | | |
|---|---|---|---|---|---|
| **Current Version** | 4 | 5 | 6 | 20 | 21 |
| 3 | 0.14 | 0.37 | 0.07 | 0.10 | 0.05 |
| 4 | . | 0.40 | 0.13 | 0.19 | 0.02 |
| 5 | . | . | 0.31 | 0.26 | 0.20 |
| 6 | . | . | . | 0.49 | 0.22 |
| 20 | . | . | . | . | 0.50 |
| All | 0.14 | 0.39 | 0.28 | 0.38 | 0.43 |

Table 3: Probability of Upgrade to Newest Version (Conditional on Upgrade Eligibility)

| | Upgrade | Upgrade |
|---|---|---|
| Distance btwn Versions | -0.0691∗∗∗ | -0.0473∗∗∗ |
| | (0.00362) | (0.00333) |
| Firm Communication Month=1 | | 0.119∗∗∗ |
| | | (0.00406) |
| Firm Communication Month=1 × Distance btwn Versions | | -0.0384∗∗∗ |
| | | (0.00182) |
| Months to Next Version: $5-8$ | 0.0326∗∗∗ | 0.00957∗∗∗ |
| | (0.00193) | (0.00166) |
| Months to Next Version: $> 8$ | 0.0432∗∗∗ | 0.0218∗∗∗ |
| | (0.00277) | (0.00249) |
| Month-Year Dummies | Yes | No |
| Calendar Month Dummies | No | Yes |
| Time Trend | No | Yes |
| Newest Version Dummies | No | Yes |
| Observations | 343,913 | 343,913 |

Additional current product and version controls, and individual fixed effects included

$* \ p < 0.05$, $** \ p < 0.01$, $*** \ p < 0.001$

A's version introduction schedule in September 2012 (see Figures 1 and 2)[18]; however, the remainder of Product A's versions were introduced on an annual basis. The remaining variation in the overall version uptake rates is a function of two factors: (1) the extent to which consumers value a given

[18] Although for the most part, each of the three products' versions were introduced on an annual schedule in the observed time-period, in September 2012 there was a shift in Product A's version release schedule that aligned it with Products B and C. As a result, consumers had a shorter time window in which to upgrade to version 6, as compared to the other versions, which partially accounts for the low uptake rate of this version.

version relative to other versions and (2) the extent to which the firm encourages the adoption of the different versions. The relative use values of the product versions are unobserved and have to be inferred from the consumers' purchase, renewal and upgrade decisions[19]. The firm's decisions that encourage upgrades, on the other hand, are at least partially observable in the data.

**Spikes in Upgrades (Firm Communication Months)** Figure 2 shows that there are some months (e.g., April 2011) when the aggregate number of upgrades within the valid subscription period spike. These spikes occur only in upgrades within the valid subscription period, as compared with the product upgrades completed together with the renewal of a subscription license. Moreover, they occur on slightly different months and to varying degrees for the different products (see Figures 13 and 14 in the Appendix Section B). Both of these empirical facts suggest that the upgrade spikes are due to an unobserved firm action rather than a demand shock. Conversations with the sponsoring firm indicate that they used notifications within the software application as well as emails to encourage upgrades. The exact details of the communication policy are currently being discussed with the sponsoring firm. For now, I include controls for the months in which the upgrade spikes are observed (one for each product version) in both the descriptive analysis and the estimation. I assume that all consumers who have a valid subscription and are using a legacy version of the product receive this notification on these months and that the notifications are not targeted further[20]. In what follows, I refer to these months as firm communication months.

**Upgrade Timing** The third pattern relates to the timing of consumers' upgrading behavior. Table 3 shows the rate at which upgrade eligible consumers upgrade to the newest version as a function of the time until the next version release. I break the time until the next version release into three groups: (1) less than five months until next version, (2) five to eight months until next version and (3) more than eight months until next version (or shortly after the release of the current version). Considering only the consumer-month observations when a given consumer is eligible to upgrade to a newer version[21], I then regress the choice to upgrade to the newest version on dummies indicating the time until the next version release (group 1, less than five months until next version, serves as a baseline).

In the regression shown in column two of Table 3, I also control for the distance between current version and newest version, month-year dummies, product, version and individual fixed effects. The coefficient on the distance between current and newest version, equal to $-0.0691$, lends further support to the existence of a non-monetary cost to upgrading that is increasing with the distance from current to newest version. Moreover, the coefficients of $0.0326$ on the "5-8 months to next version release" dummy and $0.0432$ on the "more than 8 months to next version release" dummy suggest that consumers are less likely to upgrade when the release of an even newer product version is near. This correlation is consistent with consumer behavior where upgrades are costly and con-

[19]For further discussion, see Section 6.

[20]Figure 15 in Section B of the Appendix presents evidence that the communications are not targeted based on the time until subscription expiration.

[21]Months when the consumer has a valid subscription to a legacy product version.

sumers strategically delay the upgrade close to the next version release date. Thus, the frequency and timing of upgrades allows me infer the non-monetary costs of executing an upgrade.

This same pattern holds when separately controlling for the firm communication months, as shown in column three of Table 3. Here, I control for the firm communication months and include separate calendar months dummies, new version dummies (effectively controlling for a year-long period) and a time trend to control for time. From this regression, a similar negative pattern emerges between likelihood of upgrade and the distance between the current and newest versions of the product. Moreover, firm communication months are associated with a boost to the upgrade probability, but less so for consumers who use legacy product versions. Finally, as before, less months to the release of a newer version are associated with lower probability of upgrade.

# 5 A Dynamic Model of Consumer Choice

## 5.1 Overview

My descriptive analysis reveals that legacy versions account for a large percentage of product use at any given time. To better understand this empirical fact, I further document data patterns that are consistent with state dependence in product usage as well as heterogeneity in consumer valuation of the different versions. Moreover, I document a data pattern that is consistent with consumer behavior where upgrades are costly and consumers strategically wait to upgrade to the latest version when close to the next release date. Ultimately, I am looking to understand how the firm manages its pricing mix and product design under different cost and demand conditions. I develop a dynamic model of consumer choice in order to separately quantify the effect of each of these drivers of consumer behavior and the heterogeneity therein. Estimates of these primitives of consumer behavior along with the structure of the model then allow me to evaluate consumer response to alternative firm pricing and product design strategies.

The main trade-offs I study relate to the pricing of intertemporal product versions of a particular product, rather than a firm's product assortment at a given point in time[22]. Thus, I restrict the model in some innocuous ways in order to focus on consumers' decisions related to intertemporal versions of a single product. Incorporating choices over both product assortment and intertemporal product versions is possible, but would further complicate the choice set and consumer dynamics, and contribute few substantive insights: in my particular data context, approximately 91% of consumers who pay for a subscription[23] are seen purchasing a single product. As an additional simplification, I limit the subscription length to 12-months because the overwhelming majority of paid subscriptions in my data are 12 months long (97%, see Table 1).

In my model, consumers have several distinct options. At the beginning of each month ($t$), a consumer can purchase, extend a subscription and optionally upgrade their version. His options

[22] Product line length and product version and attribute pricing has been previously studied empirically by Li (2014) and Draganska, Jain (2006), respectively.

[23] Excludes consumers who use free trial only.

may vary depending on whether he has a valid subscription to the product, whether he is using an old version of the product, etc. If the consumer has a subscription to the product in a particular month, he benefits from the use value of that product in the month of purchase as well as in the following months, until the subscription expires. As in my data, I assume that the quality of the product evolves over time via annual version releases. That is, every year, there is a newer version of the product available in the market. Both the version quality and the version release schedule enter the consumer's decision-making process.

## 5.2 Choice Set and States

To model my particular data context, I introduce three types of decisions that a consumer makes: (1) whether or not to purchase a subscription to a new product license, (2) whether or not to renew a subscription to an existing product license and (3) whether or not to upgrade a product to the newest version[24]. Letting $n$ represent the new subscription license decision, $r$ represent the renewal subscription decision and $g$ represent the upgrade decision, gives

$$\begin{aligned} n &\in \{0,1\}; \\ r &\in \{0,1\}; \\ g &\in \{0,1\}. \end{aligned} \tag{1}$$

Before discussing each of the consumer decisions in more detail, I also introduce four variables representing the different states that affect the set of choices available to the consumer: (1) remaining subscription length, $m$; (2) the product version currently in use, $e$; (3) the newest product version available in the market, $\overline{e}$; and (4) time until the next version release, $s$[25]. Then,

$$\begin{aligned} &m \in \{0,1,\ldots,12\}, \text{ where } m=0 \text{ represents an expired subscription;} \\ &e \in \{0,1,\ldots,\overline{e}\}, \text{ the consumer can use new or old product versions}^{26}; \\ &\overline{e} \in \{1,2,3\}, \text{ the firm introduces three versions of the product}^{27}; \\ &s \in \{1,2,\ldots 12\}, \text{ where } s=12 \text{ represents the new version release month.} \end{aligned} \tag{2}$$

**New subscription license purchase ($n$)** A consumer obtains the product version with the latest functionality ($\overline{e}$) when he decides to purchase a new subscription license. Moreover, after purchasing a new subscription license, the consumer is entitled to use the product for 12 months after the purchase ($m = 12$), starting on the purchase month. In this model, I restrict the version of

[24]Note that even existing consumers in my data can purchase new subscription licenses.

[25]Technically, $s$ does not affect the set of choices available to the consumer; however this state is relevant for consumer dynamics, so I introduce it here for ease of future exposition.

[26]Prior to purchasing, the consumer starts off with version value zero, thus $e = 0$ denotes that the consumer has never purchased the product before. After subscription expiration, $e$ represents the last version for which the consumer had a valid subscription.

[27]See Section 6 for a discussion of choice of three versions in the model.

the new subscription license to be the latest version of the product, which is in line with the data context documented in Section 4.

**Existing subscription license renewal ($r$)** A consumer continues to use the product in its current version ($e$), which might be distinct from the latest version ($\overline{e}$), if he renews his existing subscription license. If the license is renewed after the product expires, the consumer is entitled to use the product for 12 months after the renewal ($m = 12$), starting on the renewal month. If the license is renewed a month before expiration, the consumer is entitled to use the product for 12 months after expiration, starting on the month after expiration. There are two reasons why I distinguish between purchases and renewals. First, a major feature of the data is that I observe renewals without upgrades, which is not possible in a model with only a new subscription license purchase ($n$). Secondly, new subscription license prices and existing license renewal prices differ, as discussed in Section 3.

**Product upgrade ($g$)** A consumer can use the latest version of the product ($\overline{e}$) if he upgrades from his current version ($e$) to the latest version ($\overline{e}$). Compared to the subscription decisions, the consumer does not receive an extension to the valid subscription length upon product upgrade, as he would upon new subscription purchase or renewal. Moreover, the consumer has an opportunity to upgrade only when he has a valid subscription for an older product version ($m > 0, e < \overline{e}$). I further tailor the model to the observed data context by allowing only upgrades to the latest version (rather than any version between the current and latest versions).

Figure 4 shows a sample flow of consumer decisions. In this example, the consumer first purchases a new license (12-month subscription) for the latest version $e_1$ of the product. A few months later, the firm releases a newer version of the product ($e_2$), and because the consumer has a valid subscription to the product at that time, he becomes eligible to upgrade to the this product version. A few months after the new version is released, the consumer decides to upgrade. Note again that the upgrade decision does not change the number of months of subscription remaining on the consumer's product license. A month from the subscription expiration, the consumer renews his subscription to the product in its current version $e_2$, and the process repeats. A few months later, the consumer is again eligible to upgrade to a newer version of the product, but this time he chooses not to upgrade and, instead, lets the subscription expire.

Figure 4: Sample Consumer Purchase, Renewal and Upgrade Decision Flow

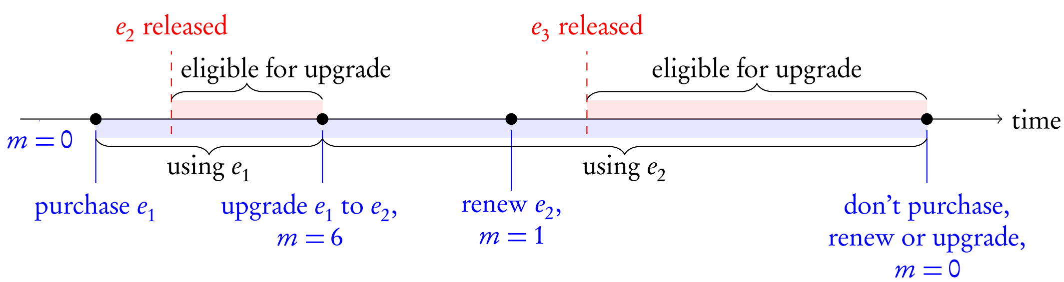

Table 4: Choice Availability,
Based on Remaining Months ($m$) and Current vs. Newest Version ($e$ vs. $\bar{e}$)

| STATES | | | | CHOICES | | | |
|---|---|---|---|---|---|---|---|
| | | Ownership | Version | Purchase | Renew | Upgrade | Renew & Upgrade |
| **Never Purchased** | 1. | $m=0$ | $e=0$ | ✓ | | | |
| **Valid Subscription** | 2. | $m>1$ | $e=\bar{e}$ | ✓ | | | |
| | 3. | $m>1$ | $e<\bar{e}$ | ✓ | | ✓ | |
| | 4. | $m=1$ | $e=\bar{e}$ | ✓ | ✓ | | |
| | 5. | $m=1$ | $e<\bar{e}$ | ✓ | ✓ | ✓ | ✓ |
| **Expired Subscription** | 6. | $m=0$ | $e=\bar{e}$ | ✓ | ✓ | | |
| | 7. | $m=0$ | $e<\bar{e}$ | ✓ | ✓ | | ✓ |

As is evident from the sample decision flow, in this context, a consumer may find himself in a number of different states. He may be using a license that is close to expiration, in which case he can choose to renew his existing license, to subscribe to the latest and greatest functionality by purchasing a new license or to let his license expire. Alternatively, he may be far from license expiration, but eligible to upgrade to the newest version of the product, in which case he can choose to upgrade or continue to use the product in its current version. Or his subscription may have expired, in which case he may be weighing whether or not to repurchase or renew the product subscription. The full choice set in the different states of the world is shown in Table 4. This table reflects most of what has already been discussed above: a renewal option is only available close to or after subscription expiration, and the upgrade option is only available when the consumer has a valid subscription to an older version. Moreover, note that, in principle, it is always possible for the consumer to purchase a new subscription license, even when he has a valid subscription for the product already[28]. Finally, the renewal and upgrade decisions may coincide, such that in one month, the consumer renews the subscription and upgrades to the latest version of the product. Although in terms of remaining subscription months and current version, this combination of decisions will result in the same outcome as the decision to purchase a new subscription license, the consumer is paying a different price and incurring different non-monetary costs when making a purchase versus a renew and upgrade decision. I allow for this flexibility in my model in order to better fit the data.

In what follows, I denote a consumer's choice by $a$ and the consumer's state by $x_t$. The state $x_t$ combines the elements discussed in the beginning of this sub-section; that is, $x_t = (m_t, e_t, \bar{e}_t, s_t)$. As shown in Table 4 and described in previous paragraphs, $a$ may include a combination of decisions, and the availability of choices $a$ differs by state. The exhaustive set of choices available to a consumer across all states is $a = (n, r, g) \in \{(1,0,0), (0,1,0), (0,0,1), (0,1,1), (0,0,0)\}$.

[28]Although in my data I observe very few consumers choosing to purchase a new license in states 2 and 3, I allow for this option in the discrete choice model in order to provide an outside option to simply using the product when far from expiration.

## 5.3 Per Period Utilities

In this sub-section I define the per period utility the consumer obtains from usage of the software product. The value to the consumer of making a purchase, renewal or upgrade choice consists of this current period utility and the discounted net present value of the future utility flows, given today's choice and the resulting evolution of the future states. I first define the current-period choice-specific utilities. In a later sub-section, I discuss the dynamic component of the value to the consumer of taking a particular action today.

Prior to discussing the choice specific per period utilities in a particular state (Equation 4), I first present the major components comprising the per period utilities. The four main components of the current period utilities are (1) product version usage utility, (2) purchase disutility, (3) renewal disutility, and (4) upgrade disutility. The product version usage utility is the utility a consumer gains from using a particular version of the product. In my model, three versions of the product are released. As a result, the version usage utility is $\theta_{1,\bar{e}} \in \{\theta_{1,1}, \theta_{1,2}, \theta_{1,3}\}$.

The second component, the purchase disutility, consists of two elements: the monetary disutility of paying a new subscription license price ($\theta_2 \cdot p_n$) and the non-monetary disutility of making the purchase ($\theta_3$). The non-monetary disutility of making the purchase can be interpreted as a transaction, hassle or search cost of seeking out an opportunity to purchase a new license, and has been used in previous research to explain low purchase incidence and low demand for variety (Hartmann, Nair 2010, Bronnenberg 2015). In my application, the new subscription license purchase cost helps explain why a consumer may choose to renew a license for an older version of the product at a higher price rather than find an opportunity to purchase a new license for the latest and greatest version of the product at a lower price (e.g., through a different channel).

The third component, the renewal disutility, consists of elements similar to the purchase disutility: the monetary disutility of paying a renewal price ($\theta_2 \cdot p_r$) and the non-monetary disutility of making the renewal ($\theta_6$). Similarly to the non-monetary disutility of making the purchase, the cost of renewal is a transaction, hassle or search cost of renewing the existing subscription license. In the case of the software application, this may consist of the consumer clicking a button that takes him to the firm's website, where the renewal payment is processed. Since there are only two types of purchases in my application, purchase of a new license and renewal, only one type of unobserved non-monetary cost is identified[29]. As a result, I normalize the non-monetary cost of renewal to 0 ($\theta_6 = 0$), and interpret the non-monetary purchase cost ($\theta_3$) as the cost of seeking out the purchase of a new license rather than renewing the existing license.

The fourth component, the upgrade disutility, consists of two elements: the non-monetary disutility of upgrading the product to the newest version of the product ($\theta_4$) and the boost to the likelihood of upgrade that comes on a firm communication month[30] ($\theta_5$). The non-monetary disutility of upgrading may include transaction or hassle costs associated with downloading and installing the newest version of the product (and resolving any incompatibility between the new product version

[29]For more discussion on parameter identification, see Section 6.

[30]Firm communication months, as discussed in Section 4, are months when consumers eligible to upgrade to the newest version are assumed to receive communications from the firm encouraging an upgrade to the latest product version.

and the consumer's operating system, for instance[31]) and learning the new product functionality. The boost in upgrading utility associated with a firm communication month reflects the extent to which a firm communication, as described in Section 4, helps the consumer to upgrade to the newest version of the product.

In addition to the per-period utility, in each period and each state, there is a shock to the consumer's utility $\varepsilon$, known to the consumer, but unobservable to the econometrician. I assume the choice-specific, independently and identically distributed (i.i.d.) extreme value error terms, $\varepsilon \sim \mathscr{EV}(\mu = 0, \sigma = 1)$, are additively separable from the other per period utility terms such that

$$u_a(x_t, e_t) = u_a(x_t) + \varepsilon_{a,t}. \tag{3}$$

To further illustrate how these components come together to form a consumer's utility from a particular action in a particular state, I consider a state when all the possible actions are available to the consumer, $x_t = (m_t, e_t, \overline{e}_t, s_t) = (1, 1, 2, 6)$. Letting $p_n$ and $p_r$ represent new license purchase and renewal prices, respectively, I consider each choice separately:

$$u_a(x_t) = \begin{cases} \theta_{1,1} & \text{if } a = (n = 0, r = 0, g = 0); \\ \theta_{1,2} + \theta_4 + \theta_5 \mathscr{I}\{t = t_{\text{comm}}\} & \text{if } a = (n = 0, r = 0, g = 1); \\ \theta_{1,1} + \theta_2 p_{rt} + \theta_6 & \text{if } a = (n = 0, r = 1, g = 0); \\ \theta_{1,2} + \theta_2 p_{rt} + \theta_6 + \theta_4 + \theta_5 \mathscr{I}\{t = t_{\text{comm}}\} & \text{if } a = (n = 0, r = 1, g = 1); \\ \theta_{1,2} + \theta_2 p_{nt} + \theta_3 + \theta_4 + \theta_5 \mathscr{I}\{t = t_{\text{comm}}\} & \text{if } a = (n = 1, r = 0, g = 0). \end{cases} \tag{4}$$

In this state, the consumer weighs five different options. He can do nothing, in which case continues to use the product in its current version 1 for one more period - the current month. He can upgrade the product to version 2, in which case, he will now get the use utility of version 2 in the current month and incur the non-monetary cost of upgrading. The cost of upgrading is offset by the firm communication boost in months where the firm takes efforts to move individuals to the newest version of the product. The consumer can also renew the product, in which case he continues to use the product in its current version and also incurs the monetary cost of renewal and the non-monetary cost of renewal. Moreover, the consumer can combine the upgrade and renewal choices, which allows him to use the latest version 2, but incurs, in an additive fashion, the renewal and upgrade costs. Finally, the consumer can choose to purchase a new subscription license for the latest version 2 of the product, which allows him to use the latest version 2, but also incurs, in an additive fashion, the purchase and upgrade costs. Note that upgrade costs are bundled in all new license purchases - this comes directly from the empirical context and my specification that the new license purchases are for the newest version of the product.

[31]Ellison and Fudenberg (2000) also consider these costs when exploring the social optimality of frequent software upgrades.

## 5.4 State Evolution

In this sub-section, I explicitly define the evolution of individual states $(m_t, e_t)$ and aggregate $(\overline{e}_t, s_t)$ states. Their law of motion drives the dynamics of $x_t = (m_t, e_t, \overline{e}_t, s_t)$. When making current-period choices, the consumer takes into account the evolution of $x_t$ and how his decisions impact the this evolution.

**Individual States** The two individual states in this context are $m$, the remaining subscription length, and $e$, the current product version. The evolution of the remaining subscription length ($m$) is governed by the new license purchase and renewal choices: the consumer gets 11 additional months of use when purchasing a new license or when renewing after the subscription has expired ($m = 0$). The consumer gets 12 additional months of use when renewing the license before expiration ($m = 1$)[32]. In the absence of purchase and renewal choices, the remaining subscription length decreases deterministically: each month there is one month less of subscription remaining. The upgrade decision does not affect the remaining subscription length. Formally, the law of motion of $m$ is given by:

$$m_{t+1} = \begin{cases} 11 & \text{if } a = (1,0,0), \text{ or } a \in \{(0,1,0),(0,1,1)\} \text{ and } m_t = 0; \\ 12 & \text{if } a \in \{(0,1,0),(0,1,1)\} \text{ and } m_t = 1; \\ \max\{m_t - 1, 0\} & \text{if } a \in \{(0,0,0),(0,0,1)\}. \end{cases} \tag{5}$$

The evolution of the product version the consumer is currently using ($e$) is governed by new license purchase and upgrade choices: the consumer upgrades to the newest functionality when purchasing a new license or when making the choice to upgrade to the newest functionality (either coupled with a renewal choice or not). The renewal decision by itself does not affect the version of the product the consumer uses. Formally, the law of motion of $e$ is given by:

$$e_{t+1} = \begin{cases} \overline{e}_t & \text{if } a \in \{(1,0,0),(0,0,1),(0,1,1)\}; \\ e_t & \text{otherwise.} \end{cases} \tag{6}$$

**Aggregate States** The two aggregate states in my context are $\overline{e}$, the latest version of the product in the market, and $s$, the time until the next version release. Both of these states are determined by time. In my model, the firm releases three intertemporal versions of the product, with consecutive version releases a year apart. Thus, the latest version is governed by the release dates and does not change between releases. Formally, the law of motion of $\overline{e}$ is given by:

$$\overline{e}_{t+1} = \begin{cases} j & \text{if } t = t_{\text{release j}}, \quad \forall j \in \{1,2,3\}; \\ \overline{e}_t & \text{if } t \neq t_{\text{release j}}, \quad \forall j \in \{1,2,3\}. \end{cases} \tag{7}$$

The time until next version release is 12 months on the month the new version is released, and

[32] This second option is a feature of my empirical context, so I allow for this transition in the model as well.

decreases by one month every month after that. Formally, the law of motion of $s$ is given by:

$$s_{t+1} = \begin{cases} 12 & \text{if } t = t_{\text{release j}}, \quad \forall j \in \{1,2,3\}; \\ s_t - 1 & \text{if } t \neq t_{\text{release j}}, \quad \forall j \in \{1,2\}. \end{cases} \tag{8}$$

After the third version is released, I assume no additional versions[33]. As a result, the time until the next version release ($s$) becomes irrelevant, and the only states that the consumer then considers are the remaining subscription months, the version of the product he is currently using, and the latest version of the product in the market ($\overline{e} = 3$).

Finally, it is important to note that although the firm communication enters the per period utility, consumers do not have any foresight regarding such communication. That is, firm communication is completely unexpected and is not taken into account in the continuation value of a particular choice[34]. As a result, these months are not a part of the consumers' aggregate state space.

## 5.5 Prices

Note that in the previous sub-section discussion, I have not specified consumers' expectations over prices. In principle, consumers have beliefs over the prices they may see tomorrow, unconditionally or conditionally on the price observed today. In my data, however, the 12 month subscription prices are reasonably level over the study period. In particular, for the specific selected product (Product A, see Section 6) in the estimation time-frame (February 2010 to August 2012), the 12 month renewal prices change in one month prior to 2012 and see a one-time permanent increase after 2012. The new purchase prices are more variable; however, in the relevant time-frame, there are long periods of constant new purchase prices, interrupted by a few changes (see Figures 9 and 10 in Appendix Section A). Given these patterns, I assume that consumers expect to encounter the prices they see in a given period in all subsequent periods. These adaptive price expectations reduce the dimensionality of my state space, and are not very restrictive in my particular data context.

## 5.6 Value Functions

The version introduction schedule and costly upgrades to the latest product version introduce a dynamic component to the consumer's choice. With positive upgrade costs (negative $\theta_4$) and increasing version valuations, the consumer may time his upgrades such that he does not have to incur the upgrade costs too frequently. Moreover, the consumer subscribes knowing that he will have the option to upgrade when the new version arrives. When a consumer is weighing a renewal or repurchase decision, he may choose to reactivate or to wait to do so, depending on when the next version will be introduced, whether it will provide better functionality than the current version and whether upgrading is costly. A consumer taking up the new version today, either through an

[33]This enables a stationary problem and likely approximates a consumer's conceptualization of the future.

[34]This assumption is made in the model because in the empirical context I observe firm communication occurring on different calendar months in different years rather than on a particular cycle.

upgrade or reactivation and upgrade decision, incurs a cost today, but changes his available options and use-value of the product tomorrow. As a result, in order to properly recover the demand primitives, I allow for forward-looking behavior on the part of the consumer with respect to the timing of version releases relative to the time remaining on the consumer's subscription.

Given the dynamic nature of my empirical context, when weighing the different purchase, renewal and upgrade choices, in the model, the consumer takes into account the current period utility he would get from a particular choice as well as the value of future choices. A forward-looking consumer is aware that the choice made today affects the state of the world and available choices tomorrow (Rust 1987). I first write a general form of the choice-specific value function, as a function of the states $x_t$, tomorrow's state-dependent choice set $\tilde{a}$, discount rate $\delta$, and choice-specific error term ${\varepsilon_t}$[35]. The extreme value i.i.d. choice-specific error term allows for a closed-form expression for the flow of future utility given the choice made today[36]:

$$\begin{aligned} v_a(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max_{\tilde{a}} \big[ v_{\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1} \big] \Big], \ \forall a \\ &= u_a(x_t) + \delta \int_\varepsilon \max_{\tilde{a}} \big[ v_{\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1} \big] dF(\overrightarrow{\varepsilon}_{t+1} | x_{t+1}), \ \forall a \\ &= u_a(x_t) + \delta \Big[ \Gamma + \log \sum_{\tilde{a}} \exp(v_{\tilde{a}}(x_{t+1})) \Big], \ \forall a. \end{aligned} \tag{9}$$

The remainder of this sub-section provides more specificity on the dynamics of consumer choice, using states 5 and 6 (as described in Table 4) as examples. I present the full set of choice-specific value functions in Appendix Section C.

**Upgrade Eligible and Close to Expiration of Existing Subscription** First, I consider a consumer whose subscription cycle coincides with the version release cycle. Let consumer's state $x_t = (m_t, e_t, \overline{e}_t, s_t) = (1,1,2,1)$; that is, the legacy consumer is a month away from subscription expiration of his product version 1, and version 3 is scheduled to be released in the following month. In this state, the consumer's choice set consists of all the possible choices: do nothing, purchase a new subscription license, renew the subscription, upgrade and renew and upgrade. Using the states described in Table 4 as indices, I write the consumer's choice-specific value functions in this particular state as follows:

$$\begin{aligned} v_{5,a}(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max \big[ v_{3,\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1} \big] \Big], \forall a \in \{(1,0,0),(0,1,1),(0,1,0)\}; \\ &\qquad \forall \tilde{a} \in \{(1,0,0),(0,0,1),(0,0,0)\}; \\ v_{5,a}(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max \big[ v_{7,\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1} \big] \Big], \forall a \in \{(0,0,1),(0,0,0)\}; \\ &\qquad \forall \tilde{a} \in \{(1,0,0),(0,1,1),(0,1,0),(0,0,0)\}. \end{aligned} \tag{10}$$

[35] In my application, the individual and aggregate states evolve deterministically. As a result, the uncertainty is only over the future draws of the error terms.

[36] $\Gamma$ represents the Euler's constant.

The choice to purchase a new subscription license, renew an existing license and renew and upgrade the product all result in a similar state in the following period: a state where the consumer has multiple months of a valid subscription left and is eligible to upgrade to a newer version of a product. Separately, the choice to upgrade to the new version and the choice to do nothing also result in similar state in the following period: a state where the consumer's subscription has expired, and he is eligible to purchase a new subscription, renew an existing subscription or renew the subscription together with a product upgrade.

If the consumer chooses to purchase a new subscription license in the current period, in the following period, he will have 11 more months remaining on his subscription; however, since a new version will come out in the following period, in the following period he will be using an old version of the product and weighing the decision to upgrade to the newest version or not. By purchasing a new license in the current period, the consumer incurs an upgrade cost today and faces the possibility of incurring the upgrade cost tomorrow. In addition, by purchasing in the current period, the consumer loses a month of subscription, since the subscription starts on the month of new license purchase. The state evolution from the renewal and upgrade choice is similar, except if the consumer chooses to renew and upgrade, he will have 12 more months remaining on his subscription in the following period[37]. On the other hand, if the consumer chooses to renew the subscription in its current version only, he will not incur upgrade costs in the current period, but will face the same upgrade costs tomorrow that he is facing today. As a result of these subscription and upgrade dynamics, with increasing valuations for successive versions and costly upgrades (negative $\theta_4$), the consumer may prefer the renewal choice to the new license purchase or renewal and upgrade choice in this state.

If the consumer chooses to do nothing, his license expires. In the following month, however, in order to continue subscription coverage, the consumer can either purchase a new subscription license, renew the expired subscription in its old version or renew the subscription together with a product upgrade to the latest version. Moreover, as with the choice to renew, when choosing to wait, the consumer does not incur upgrade costs today only to face them again tomorrow. On the other hand, if the consumer chooses to upgrade the product to the latest version in the current period, he will face upgrade costs again in the following period when a newer version of the product becomes available. With increasing valuations for successive versions and costly upgrades (negative $\theta_4$), the consumer may prefer to wait over upgrading to the newer version of the product in the current period.

A consumer whose subscription cycle does not coincide with the version release cycle faces similar trade-offs, but, depending on the time until the next version release and his successive version valuation, this consumer may have some incentive to upgrade to the latest version of the product in

[37] This is because a renewal action before expiration adds 12 subscription months to the expiration date. Intuitively, the consumer may want to renew the subscription a month prior to expiration because of some additional convenience that may exist this month, but not in the following month. For instance, if the firm sends a renewal reminder a month prior to expiration, the consumer may want to take the renewal action today rather than waiting until next period. In my model, the current period error shock relative to the expectation over future error shocks will account for this kind of unobserved convenience of renewing in a particular month.

the current period. Let the consumer's state $x_t = (m_t, e_t, \overline{e}_t, s_t) = (1,1,2,8)$; that is, the consumer is a month away from subscription expiration of his product version 1, and version 3 is scheduled to be released 8 months from now. Using the states described in Table 4 as indices, I write the consumer's choice-specific value functions in this particular state as follows:

$$\begin{aligned}
v_{5,a}(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[\max\big[v_{2,\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1}\big]\Big], \forall a \in \{(1,0,0),(0,1,1)\}; \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall \tilde{a} \in \{(1,0,0),(0,0,0)\}; \\
v_{5,a}(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[\max\big[v_{3,\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1}\big]\Big], \text{ if } a = (0,1,0); \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall \tilde{a} \in \{(1,0,0),(0,0,1),(0,0,0)\}; \\
v_{5,a}(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[\max\big[v_{6,\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1}\big]\Big], \text{ if } a = (0,0,1); \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall \tilde{a} \in \{(1,0,0),(0,1,0),(0,0,0)\}; \\
v_{5,a}(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[\max\big[v_{7,\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1}\big]\Big], \text{ if } a = (0,0,0); \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall \tilde{a} \in \{(1,0,0),(0,1,1),(0,1,0),(0,0,0)\}.
\end{aligned} \tag{11}$$

In this state, the main additional dynamics to consider for a consumer whose renewal cycle is different than the version release cycle are around the choice to upgrade. If the consumer chooses to purchase a new license or renew the subscription and upgrade to the latest version, he will be able to use the product in version 2 for another 8 months before version 3 is introduced. Alternatively, if the consumer chooses to renew the product only, he will continue to be eligible for an upgrade tomorrow. The consumer may not want to purchase a new license today, since that would mean losing a month of subscription. Whether the consumer decides to upgrade (paired with renewal or not), however, will depend on the size of his non-monetary upgrade costs and how these costs compare to the incremental value the consumer places on version 2 over version 1. For instance, the consumer may want to renew the product without upgrading[38] if he gets positive net present value from using version 1 for the next 12 months, and the net present value of the difference between his version 2 and version 1 use values is less than the upgrade costs he incurs today.

On the other hand, whether the consumer decides to re-purchase or renew at all will depend on the monetary and non-monetary costs of renewal and purchase as compared to the value he gets from using the product. For instance, if his net present value from using version 1 is lower than the monetary and non-monetary renewal or purchase costs, and the upgrade costs are high enough, the consumer may allow the subscription to expire and never re-purchase or renew the product in any version.

**Expired Subscription and Previous Usage of Latest Version** Now let the consumer's state $x_t = (m_t, e_t, \overline{e}_t, s_t) = (0,2,2,8)$; that is the consumer last used version 2 of the product, but

[38]Or wait to renew without upgrading until the following period.

the subscription is currently expired. In this case, each period, the consumer is deciding whether or not to purchase a new subscription license or renew the existing license. Using the states described in Table 4 as indices, I write the consumer's choice-specific value functions in this particular state as follows:

$$
\begin{aligned}
v_{6,a}(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[\max\big[v_{2,\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1}\big]\Big], \forall a \in \{(1,0,0),(0,1,0)\}; \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad \forall \tilde{a} \in \{(1,0,0),(0,0,0)\}; \\
v_{6,a}(x_t) &= u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[\max\big[v_{6,\tilde{a}}(x_{t+1}) + \varepsilon_{\tilde{a},t+1}\big]\Big], \text{ if } a = (0,0,0); \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad \forall \tilde{a} \in \{(1,0,0),(0,1,0),(0,0,0)\}.
\end{aligned}
\tag{12}
$$

Thus, if the consumer purchases in the current period, he is covered for an additional 11 months in the following period. If he chooses to do nothing in the current period, however, the consumer weighs the same choices in the following period. One reason a consumer may choose to wait is that he is waiting for the newer version, version 3, to come out before renewing or re-purchasing. This would be the case if the consumer does not value version 2 very much, but does value version 3 significantly more. Additionally, the consumer may choose to purchase in a particular period because the purchase price is low: in this case, the net present value of the consumer's purchase utility would be higher than in a period with a higher purchase price. Finally, the consumer may choose to purchase in a particular period because his error draw in that period is particularly high.

By discussing the choice-specific value functions in these two states, I have covered the main consumer dynamics driving my results. As a result of the firm's subscription and pricing policies, the consumer may find himself in a number of other states of the world. I present the value functions for each of these states in the Appendix Section C, but do not discuss them in further detail here.

## 5.7 Computation

For a given set of parameters $(\theta_{1,1}, \theta_{1,2}, \theta_{1,3}, \theta_2, \theta_3, \theta_4)$[39], I compute each of the finite horizon choice-specific value functions $v_a$, using $T = 100$ months[40]. Effectively, this means that after version 3 comes out, the consumer knows that for another 100 months, he will have the option to purchase, renew and upgrade to version 3 of the product, but no additional versions of the product will be released. In my data, I observe six additional months after version 3 is released. Thus, in the last month of my data set, a consumer weighing the different choices in his particular state will have a 93 month continuation value for each of his choices. I then compute the continuation value for

[39]Since the consumers do not anticipate the firm communication shocks, $\theta_5$ does not enter the value function computation and affects the current period utility only. Moreover, $\theta_6 = 0$. As a result, both $\theta_5$ and $\theta_6$ are omitted from the above set of parameters.

[40]I choose $T$ to be sufficiently high in order to minimize the impact of the terminal zero value on the shape of the consumer's value function and his current-period behavior. For instance, the net present value of 1 util received 100 months from now, at a discount rate of $\delta = 0.95$ is 0.0059. Thus, the numerical difference between $T = 100$ and $T = \infty$ is small, allowing me to approximate $T = \infty$ with less computational burden.

choices in previous periods from the continuation values during the periods when version 3 is the latest version ($\overline{e} = 3$).

# 6 Estimation

## 6.1 Estimation with Unobserved Heterogeneity

The model specification developed in the previous section specifies a set of parameters $\Theta = (\theta_{1,1}, \theta_{1,2}, \theta_{1,3}, \theta_2, \theta_3, \theta_4, \theta_5)$ as the primitives driving consumer demand. The importance of heterogeneity in consumer behavior is well documented in the literature and suggested by some of the patterns illustrated in Section 4 here. To simplify computation, I allow for a two-type discrete mixture in the demand parameters, where $\alpha$ is the incidence of type A consumers and $1-\alpha$ is the incidence of type B consumers:

$$\Theta_w = \begin{cases} \Theta_A & \text{with probability } \alpha; \\ \Theta_B & \text{with probability } 1-\alpha. \end{cases} \tag{13}$$

I estimate the parameter vectors $\Theta_A, \Theta_B$ and $\alpha$ using maximum likelihood. Given the assumption of choice-specific, identically distributed extreme value error terms, the probability of observing an individual $i$'s action $a$ in a state $x_t$ with prices $(p_{nt}, p_{rt})$ under a set of parameters $\Theta_w$ is[41]

$$\mathscr{P}\{a_{it}; x_{it}, p_{nt}, p_{rt}, \Theta_w\} = \frac{\exp(v_a(x_{it}, \Theta_w))}{\sum_{\tilde{a}} \exp(v_{\tilde{a}}(x_{it}, \Theta_w))}. \tag{14}$$

The likelihood of the data is then

$$\mathscr{L}(\Theta) = \prod_{i=1}^{I} \prod_{t=1}^{T_i} \left[ \alpha \mathscr{P}\{a_{it}; x_{it}, p_{nt}, p_{rt}, \Theta_A\} + (1-\alpha) \mathscr{P}\{a_{it}; x_{it}, p_{nt}, p_{rt}, \Theta_B\} \right]. \tag{15}$$

As discussed above, the estimated set of parameters is $\Theta = (\Theta_A, \Theta_B, \alpha)$. The discount factor ($\delta$) cannot be separately identified in this empirical context. As a result, I fix the discount factor to $\delta = 0.95$ and estimate the remainder of the parameters given this value of the discount rate[42].

## 6.2 Estimation Procedure

I estimate the parameters $(\Theta_A, \Theta_B, \alpha)$ in the following manner: first, for an arbitrary guess for the set of parameters $(\Theta_A, \Theta_B, \alpha)$, I compute current period utilities and continuation values for each action in each state, as described in Section 5, separately for each type. Given these type-specific values and the parameter $\alpha$, I then compute the data likelihood. Based on this value, the

[41] $\tilde{a}$ denotes all actions that are accessible to the consumer in state $x_{it}$.

[42] I also test discount rates of $\delta = 0.98$ and $\delta = 0.9$. While the choice of discount rate affects the exact point estimates, the estimation results and conclusions are directionally the same as with a $\delta = 0.95$ discount factor.

optimization routine then takes another parameter guess, and the process repeats itself and iterates until convergence[43].

## 6.3 Estimation Data and Assumptions

In order to ensure that the model fits well with my empirical application, when specifying the model, I account for the vast majority of factors specific to the product design and pricing of this particular firm (e.g., the availability of purchase and renewal options, the different possible states, etc.). There are a number of smaller features of the product design and pricing that are specific to the firm, but are not fundamental to my research question. To avoid further complicating the state space and choices, I do not account for them in my model specification. Moreover, for the empirical application, I choose the product that is most consistent with my model assumptions.

I present the comprehensive list of the data selection criteria in the Appendix Section D. Here, I briefly discuss the reasons behind choosing Product A and the time period between February 2010 and August 2012 for estimation. First, of the three products, Product A has the highest number of paid subscriptions. Product A has the second highest percentage of observations retained after applying the data selection criteria (75% for Product A, 80% for Product B); however, in comparison to Product B, there are about twice as many purchases (new subscription license purchases and renewals) and about four times as many upgrades for Product A. Product A also has the highest re-purchase and renewal rate[44] of the three products. A rich panel dimension is important in identifying unobserved heterogeneity in consumer behavior. Moreover, in order to separately identify the non-monetary purchase and upgrade costs, multiple observations of purchases and upgrades are required. As a result, of the three products that the firm offers, I use Product A for estimation.

Furthermore, my model captures the introduction of three versions, but during the time-span when price data is available, six distinct versions of the product are the "newest" available in the market. In principle, increasing the number of distinct intertemporal versions may help with conclusions regarding the effectiveness firm's long-term innovation strategy, however, it also complicates computation. The main dynamics and usage utility trends are captured with three product versions in the model, and to estimate these parameters, I keep data for three intertemporal versions of Product A only (version 4, 5 and 6, re-labeled to version 1, 2 and 3, respectively). In particular, I focus on the time period before the shift in version release schedule on September 2012. I keep only those consumers who first purchase on or after February 2010, when version 4 comes out, and I truncate the panel after August 2012, when version 20 comes out. After applying these additional selection criteria, I arrive at a dataset containing $5,078$ consumers of product A and the following choices

[43] In Appendix Section E, I present a simulation exercise for a homogeneous model to illustrate that this procedure is well-suited to recover the parameters of interest.

[44] Re-purchases and renewals are purchases that the consumer makes after having already purchased once in the past.

made by these consumers:

$$(n,r,g) = \begin{cases} (1,0,0): \text{New Subscription License Purchase} - 5,394^{45}; \\ (0,1,1): \text{Existing Subscription License Renewal and Product Upgrade} - 106; \\ (0,1,0): \text{Existing Subscription License Renewal} - 1,098; \\ (0,0,1): \text{Product Upgrade} - 942; \\ (0,0,0): \text{No Action} - 69,878^{46}. \end{cases} \tag{16}$$

## 6.4 Identification

In this sub-section, I present an informal discussion of parameter identification. The non-monetary cost of making the upgrade ($\theta_4$) is identified from the rate of upgrade conditional on upgrade eligibility; i.e., how long the consumer chooses to wait before upgrading in states of the world when he is eligible to do so: the longer the wait, the higher the upgrade costs. Moreover, the upgrade costs are identified relative to the the discount rate, since the discount rate determines the extent to which the possibility of incurring upgrade costs in the future matter to the consumer today. The benefit of a firm communication month ($\theta_5$) is identified from the extent to which upgrade eligible consumers are more likely to upgrade on firm communication months.

The non-monetary cost of purchasing a new subscription license ($\theta_3$) is identified from the consumer's choice of a new subscription license purchase rather than renewal of an existing subscription. Since the renewal choice is simpler than a re-purchase choice, I normalize the renewal cost to 0 and estimate the purchase cost relative to the renewal cost. Similarly to the upgrade costs, the non-monetary purchase cost is identified relative to the discount rate. That is, the extent to which a consumer cares about incurring costs of purchasing in the future affects whether he chooses to purchase today.

The monetary cost of purchasing a new subscription or renewing an existing license ($\theta_2$) is identified from joint intertemporal and cross-sectional price variation. The price paid by consumers varies somewhat over time (Figures 9 and 10), and the extent to which consumers are more likely to reactivate the subscription at lower prices will inform their price sensitivity. Moreover, the price paid also varies between a new subscription license purchase and renewal of an existing license. As a result, the price coefficient is jointly identified from the choice of one of these options over another, as the difference in the prices varies over time, and the timing of the choice.

The use utility parameters ($\theta_{1,1}, \theta_{1,2}, \theta_{1,3}$) are empirically identified from the differences in rates of upgrade, purchase and renewal of different versions (see Figure 3). Because these utility parameters are identified from rates at which consumers take specific actions, differences in these parameters reflect the consumers' perceptions of the improvements made between versions as well as any differences in costs of adopting different versions. These parameters are also identified relative to

[45] 316 are re-purchases, 5,078 are first-time purchases.

[46] $\sim$ 75% are while consumers have a valid subscription for the product, 25% are while consumers' subscriptions are expired.

the discount factor, since a purchase of the product today allows the consumer to use the product for an additional 11 or 12 months.

The heterogeneity in consumer behavior is identified from the panel dimension of the data, and the repeated purchase, renewal and upgrade decisions made by consumers.

One other important consideration is a consumer's first purchase of subscription for the product. In its current state, the model predicts that, if the use value is increasing from one version to the next, then a given consumer who has never purchased before may have incentive to wait until later versions of the product are released before ever purchasing the product. The concern is that, in addition to product quality, the arrival of a consumer's first purchase is in large part driven by the acquisition of a machine. In my data, I do not directly observe when the consumer acquired the machine on which the software product was eventually installed.

It may be possible to supplement the current consumer-level model with an aggregate model of the arrival of consumers' first purchase opportunities. In the current dataset, however, I observe roughly equally sized cohorts of consumers that represent a random sample from the firm's total population of users[47]. Estimation of an aggregate model would require supplementary data for the total number of consumers purchasing the product, the number of machines sold in the US, and the proportion of the market held by Product A in each month in my estimation period. In lieu of this data, it may also be possible to infer, albeit imperfectly, how long the consumer had the machine prior to purchasing from his use of free trials for Product A or his purchase activity and / or use of free trials for the other two products. In the current version of the paper, I take a third route and assume that machine acquisition happens when the first purchase of the product is observed. As a result, the first new subscription license of each consumer is not used for estimation. In future work, I plan to relax this assumption by accounting for the arrival of consumers' first purchases in one of the ways described above.

# 7 Estimation Results and Discussion

## 7.1 Parameter Estimates and Individual Consumer Behavior

In Table 5, I present homogeneous and heterogeneous model estimation results[48]. The homogeneous model estimation results suggest that, on average, consumers value each successive version of the product slightly more than the previous one. Consumers also have a negative utility associated with paying a price as well as negative non-monetary utility of purchasing a new subscription license (relative to renewing) and negative non-monetary utility of upgrading to the newest product version. With these parameter estimates, the net present value of using the product in the third version ($\overline{e} = 3$) for 12 months, given a discount rate of $\delta = 0.95$, is 9.8262. The disutilities of a new purchase and renewal at the mode prices ($p_n = 6$, $p_r = 8$), on the other hand, are $-10.947$ and $-11.466$, respectively. Even at the lowest observed new purchase and renewal prices ($p_n = 4.8$,

[47]Sample percentage for each of the cohorts is not specified.

[48]The homogeneous model assumes one type, so $\Theta_A = \Theta_B$.

$p_r = 7$), the consumer's disutilities of purchase are $-9.227$ and $-10.033$, respectively. These parameter estimates suggest that consumers will typically choose not to repurchase or renew the product. Moreover, the homogeneous model estimates suggest that the upgrade costs are more negative than new subscription license purchase costs. That is, on average, consumers are less averse to seeking out an opportunity to purchase a new license at a lower price than they are to having to download and install the newest version of the software on their machine. Finally, a comparison of the price coefficient and the firm communication coefficient reveals that the boost to the consumer utility during firm communication months is equivalent to about \$0.9 off the price paid (15% and 11% of the mode purchase and renewal prices, respectively).

The homogeneous model estimates provide a benchmark for the heterogeneous results. In my application, heterogeneity in consumers' product version valuation, willingness-to-pay, etc. is especially important, since it directly impacts the firm's optimal design and pricing of product versions. In columns 4 and 5 of Table 5, I present heterogeneous model estimates, and for the remainder of the section I discuss the implications of these results.

The estimates with two type discrete mixture heterogeneity in the parameters suggest that one of the consumer types, Type A, is a price sensitive / low value consumer who, nonetheless, assigns slightly increasing value to successive versions of the product. This consumer type makes up about 3/4 of the consumer population. He has a relatively low cost of purchasing a new subscription license, as compared to renewing the existing license. In fact, eliminating the non-monetary component of the purchase cost would be equivalent to about a \$0.6 discount off the purchase price (10% of the mode purchase price), in terms of the improvement to consumer utility.

For a simple example, consider the time-frame after the release of the third and final version. A Type A consumer who is using an up-to-date version of the product is deciding only whether to repurchase or renew the subscription. His net present value of using the product in the third version ($\overline{e} = 3$) for 12 months, given a discount rate of $\delta = 0.95$, is 14.214. On the other hand, his disutilities of purchasing a new subscription and renewing an existing subscription are $-16.43$ and $-19.944$, respectively, at mode prices ($p_n = 6$, $p_r = 8$) and $-13.439$ and $-17.451$, respectively, at the lowest observed prices ($p_n = 4.8$, $p_r = 7$). Thus, a Type A consumer may repurchase or renew in months with low prices or in months when he has some other high unobserved taste shock associated with making a purchase or renewal, but he is likely to choose not to repurchase or renew at mode prices[49].

Moreover, a Type A consumer has high upgrade costs, which means that he will rarely upgrade to the newest product version, conditional on being eligible to do so, even though his valuation for the successive versions is slightly increasing, as evidenced by repurchasing behavior. Moreover, for this consumer type, the firm communication gives a relatively low boost to the upgrade utility and does little to overcome the cost of upgrading. I put this effect in dollar terms by comparing $\hat{\theta}_5$ to the price coefficient: for Type A consumers, the boost to consumer utility of upgrading during firm

[49]This behavior is further illustrated in the left panel of Figure 19 in the Appendix. Under mode prices, the consumer will not repurchase or renew the subscription, even after it expires, unless he receives a significant utility shock in that month. He is more likely to repurchase or renew version 3 over version 2, however, which is further evidenced by comparing the left panel of Figures 19 and 21 in the Appendix.

Table 5: Estimation Results

| | Homogeneous Model | Heterogeneous Model | |
|---|---|---|---|
| | | **Type A** | **Type B** |
| **Version Preferences** | | | |
| $\theta_{1,1}$ | 0.8231 | 1.4166 | 3.9348 |
| | (0.0405) | (0.0981) | (0.1991) |
| $\theta_{1,2}$ | 1.0010 | 1.5378 | 3.7906 |
| | (0.0420) | (0.1035) | (0.1995) |
| $\theta_{1,3}$ | 1.0689 | 1.5462 | 3.6699 |
| | (0.0431) | (0.1079) | (0.2021) |
| **Price** | | | |
| $\theta_2$ | -1.4333 | -2.4293 | -1.6925 |
| | (0.0484) | (0.1235) | (0.0931) |
| **Purchase Costs** | | | |
| $\theta_3$ | -2.3472 | -1.4722 | -3.8587 |
| | (0.0927) | (0.2159) | (0.1812) |
| **Upgrade Costs** | | | |
| $\theta_4$ | -4.8446 | -5.0982 | 0.0930 |
| | (0.0912) | (0.1912) | (0.3096) |
| $\theta_5$ | 1.2634 | 0.8509 | 1.5679 |
| | (0.0631) | (0.1534) | (0.0796) |
| **Type Proportion** | | | |
| $\alpha$ | — | 0.7342 | 0.2658 |
| | | (0.0447) | |

Asymptotic standard errors in parentheses

communication months is equivalent to a \$0.4 discount on the purchase or renewal price (7% and 5% of the mode purchase and renewal prices, respectively). Given the significant costs of upgrading and the increasing valuation of successive versions, Type A consumer exhibits behavior consistent with the pattern documented in Table 3. That is, Type A consumer will strategically delay upgrading when closer to the next version release, and upgrade with higher probability when a new version has just been released; however, in order for a Type A consumer to upgrade, he will need to experience a high shock for that particular choice in a given month[50].

The second consumer type, Type B, represents about 1/4 of the consumer population. This consumer is a price insensitive / high value consumer, however, his valuation over the successive versions of the product is decreasing. The homogeneous model estimates suggested an overall increase in successive version valuation; however, the heterogeneous model reveals that although this is the case for the majority of the consumers, for a high value consumer, the pattern is the oppo-

[50]This behavior is further illustrated by comparison of the two panels in Figure 23 in the Appendix.

site. The high value consumer also has a relatively higher non-monetary cost of purchasing a new subscription license, as compared to renewing an existing license. For this consumer, eliminating the non-monetary cost of seeking out an opportunity to purchase a new license rather than renew the existing one, in utility terms, is equivalent to a significant discount on the purchase price: $2.3 or 38% of the mode purchase price. Returning to the example of the consumer's choices after the release of version 3, note that, even at mode prices, consumer B is better off repurchasing and renewing the product after expiration. His net present value of using the product in the third version ($\overline{e} = 3$) for 12 months, given a discount rate of $\delta = 0.95$, is 33.737, and his disutility from purchasing and renewing at mode prices is $-14.014$ and $-13.54$, respectively. Thus, even at relatively high prices, the consumer is better off re-purchasing or renewing rather than allowing his subscription to lapse[51].

A Type B consumer also differs from the Type A consumer in terms of upgrading behavior: his valuation of successive versions of the product is decreasing; however, his non-monetary costs of upgrading to a newer version of the product are not significantly different from zero. Moreover, Type B consumer gets a significant boost to his probability of upgrading during firm communication months. In dollar terms, the boost to Type B consumer utility of upgrading during firm communication months is equivalent to a $0.9 discount on price (15% and 11% of new purchase and renewal prices, respectively). As a result, a Type B consumer is much more likely to upgrade during communication months and not likely to upgrade otherwise, given his decreasing valuation of new versions.

## 7.2 Aggregate Demand Implications

It is important to note that the individual consumer behavior implied by these estimates is fundamentally different from the consumer behavior in the typical durable goods example (e.g., cameras) described in the opening paragraph. While in the typical durable goods case, the high value, high willingness-to-pay consumers value the innovation more than the low value, low willingness-to-pay consumers, in my application this is not the case. The high value consumers do not value the innovation of the newer versions, while the low value consumers, more typically, repurchase whenever they see a sufficiently low price or have another sufficiently high unobserved shock to their purchase utility. Thus, the high value consumers continue to repurchase and renew older versions of the product, while the low value consumers repurchase rarely, but value newer versions more than older ones when they do.

The individual consumer behavior described above leads to the aggregate demand behavior illustrated in Figures 5 and 6. In the top panel of Figure 5, I plot the aggregate first subscription license

[51] This behavior is further illustrated in the left panel of Figure 20 in the Appendix. Under mode prices, the consumer does not find it beneficial to repurchase or renew the subscription prior to expiration without an additional positive shock to the utility of that choice; however, after the product expires, the consumer is likely to find it beneficial to renew or repurchase. Moreover, since a Type B consumer values newer versions less than older versions, he is less likely to repurchase version 3 than he is to repurchase version 2. This can be seen by comparing the left panel of Figure 20 to Figure 22 in the Appendix.

Figure 5: First License Purchases, by Version, by Type

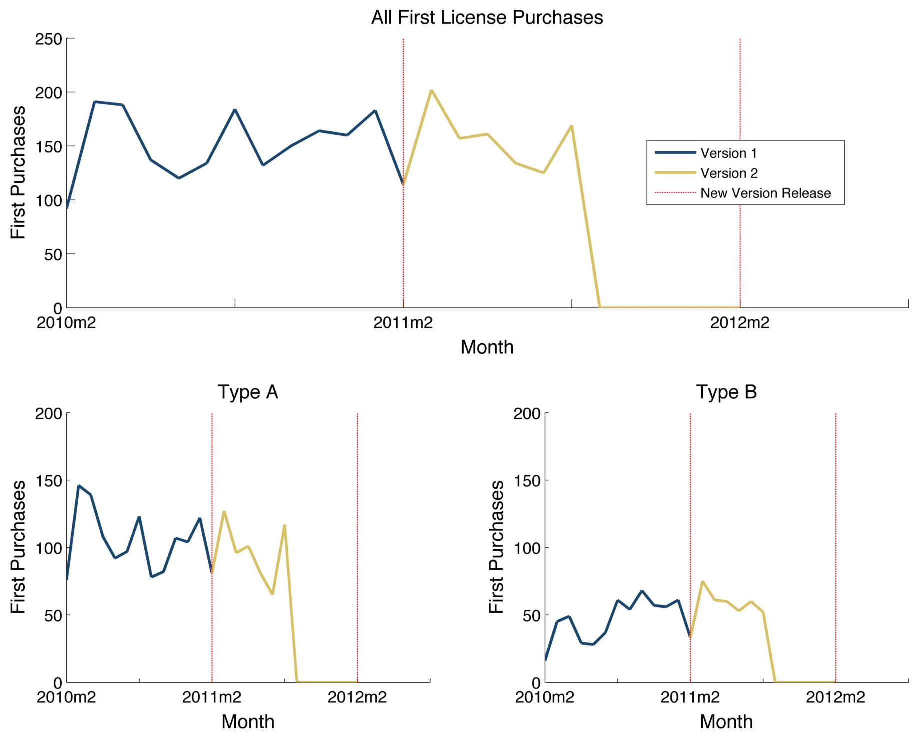


Figure 6: Re-Purchases and Renewals, by Version, by Type

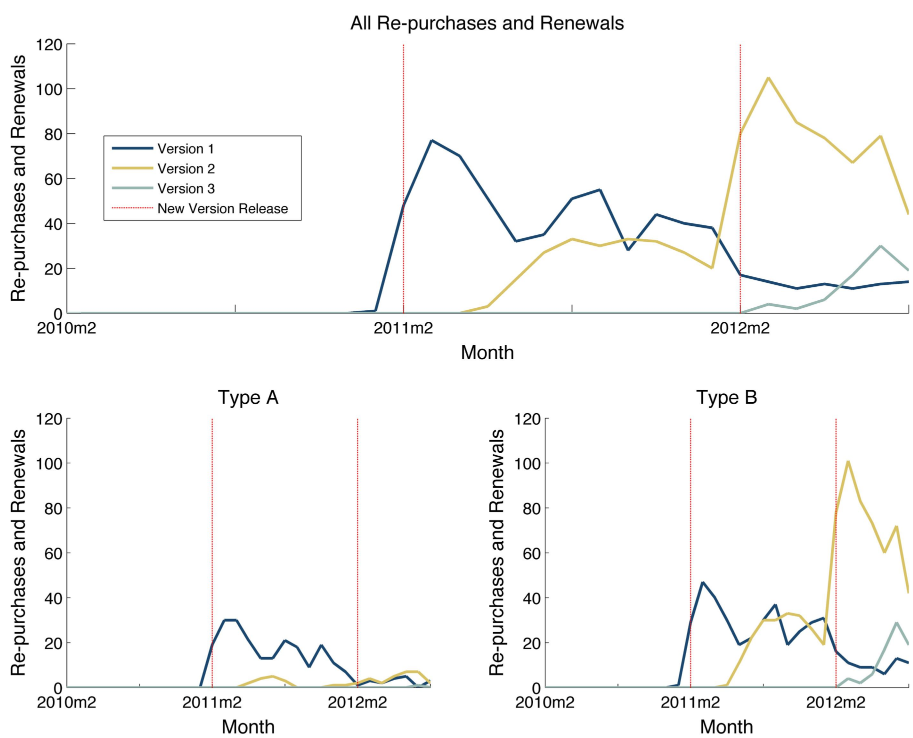

purchases of all consumers who first purchase prior to September 2011[52]. Based on the parameter estimates and individual consumer behavior, I calculate each consumer's posterior probability of being a Type A consumer. I assign a particular customer to Type A if his probability of being Type A is higher than 0.5 and to type B otherwise[53]. In the bottom two panels of Figure 5, I then plot the aggregate first subscription license purchases for the two consumer types separately. Since the Type A consumer is more prevalent, type A consumers are behind the majority of the first subscription license purchases. By examining the re-purchase and renewal behavior in Figure 6, we see that although first purchases mostly originate from Type A consumers, Type B consumers drive re-purchases and renewals. Moreover, considering re-purchase and renewals by version, we see that Type B consumers drive legacy version usage by continuing to renew older versions.

In fact, this insight may help explain why we observe differentiated new purchase and renewal prices in this context. As discussed in previous sections, when a consumer renews the existing license, he renews the subscription for the current version of the product, which may be the latest version or a legacy version. On the other hand, when a consumer purchases a new subscription license, he purchases the latest functionality. A low new license purchase price may encourage new consumers to purchase. In addition, it serves, at least partially, to sort consumers: if the high value consumers want to use an older version they know and like, they need to pay a higher price for renewing the product in its existing version without upgrading. On the other hand, low value consumers value a low price and, thus, purchase the cheapest option if purchasing at all.

In the Appendix Section H, I present a simple model of the firm's pricing decisions across versions when consumers have heterogeneous valuations for new versions of the product. In the model, it is optimal for the firm to offer new functionality at a discount, rather than pricing it at a premium, when faced with a high value consumer minority who do not value innovation and a low value consumer majority who do. Thus, the model's predictions are consistent with the observed firm behavior discussed above.

## 7.3 Purchase and Renewal Sensitivity to Price

Before evaluating alternative pricing schemes that a firm may want to implement, I first examine a temporary price change in order to illustrate the different factors affecting consumers' response to price. I use the structure of the model and the estimates to simulate a consumer's purchase, renewal and upgrade paths, conditional on their first purchase[54]. To better understand how consumers respond to price changes in this dynamic context, I then simulate a price change in a particular month (March 2011) and evaluate how consumer purchases, renewals and upgrades respond to this temporary change, relative to the simulated baseline. From the consumer's perspective, the change in price is unexpected and only enters his current period utility. That is, on the month of the price

[52]Because the subscriptions of this subset of consumers expire prior to the end of the observation period, these consumers have re-purchase and renewal opportunities observed in the data.

[53]The vast majority of this subset of consumers have either a very high ($> 0.9$) or a very low ($< 0.1$) posterior probability of being type A.

[54]The full results on the heterogeneous model fit are available in Section F of the Appendix.

change, the consumer anticipates to see one set of prices $(p_n, p_r) = (6{,}7)$, but instead faces a different set of prices $(p_n - 1, p_r)$,$(p_n, p_r - 1)$ or $(p_n - 1, p_r - 1)$. The consumer knows, however, that the change in price is temporary; i.e., he anticipates that tomorrow he will be faced with the same set of prices he expected to see today. Table 6 shows cumulative demand response to this type of price change in renewal price, purchase price or both at 1, 2, 12 and 18 months after the change.

Table 6: Cumulative Demand Response to One Time Price Change in March 2011

| $\Delta_{p_n}$ / $\Delta_{p_r}$ | Mar '11 | | | Mar-Apr '11 | | | Mar '11-Feb '12 | | | Mar '11-Aug '12 | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | $n$ | $r$ | $g$ | $n$ | $r$ | $g$ | $n$ | $r$ | $g$ | $n$ | $r$ | $g$ |
| **-\$1 / -\$1** | | | | | | | | | | | | |
| Total | 3 | 62 | 2 | 3 | 33 | 0 | -2 | 26 | 5 | -2 | 26 | 2 |
| | N/A | (117%) | (4%) | (150%) | (30%) | (0%) | (-1%) | (4%) | (1%) | (-1%) | (2%) | (0%) |
| Type A | 3 | 32 | 0 | 3 | 32 | 0 | -2 | 26 | 5 | -2 | 26 | 3 |
| | N/A | (640%) | (0%) | (300%) | (400%) | (0%) | (-2%) | (23%) | (3%) | (-1%) | (18%) | (1%) |
| Type B | 0 | 30 | 2 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | -1 |
| | N/A | (63%) | (6%) | (0%) | (1%) | (0%) | (0%) | (0%) | (0%) | (0%) | (0%) | (0%) |
| **-\$1 / —** | | | | | | | | | | | | |
| Total | 9 | -5 | 0 | 9 | -5 | 0 | 8 | -5 | -1 | 8 | -6 | -1 |
| | N/A | (-9%) | (0%) | (450%) | (-5%) | (0%) | (4%) | (-1%) | (0%) | (3%) | (0%) | (0%) |
| Type A | 4 | 0 | 0 | 4 | 0 | 0 | 3 | -1 | 0 | 3 | -1 | 0 |
| | N/A | (0%) | (0%) | (400%) | (0%) | (0%) | (3%) | (-1%) | (0%) | (2%) | (-1%) | (0%) |
| Type B | 5 | -5 | 0 | 5 | -5 | 0 | 5 | -4 | -1 | 5 | -5 | -1 |
| | N/A | (-10%) | (0%) | (500%) | (-5%) | (0%) | (6%) | (-1%) | (0%) | (4%) | (0%) | (0%) |
| **— / -\$1** | | | | | | | | | | | | |
| Total | 0 | 63 | 2 | 0 | 34 | 0 | -4 | 27 | 6 | -4 | 27 | 3 |
| | N/A | (119%) | (4%) | (0%) | (31%) | (0%) | (-2%) | (4%) | (1%) | (-1%) | (2%) | (0%) |
| Type A | 0 | 33 | 0 | 0 | 33 | 0 | -4 | 27 | 6 | -4 | 27 | 4 |
| | N/A | (660%) | (0%) | (0%) | (413%) | (0%) | (-4%) | (24%) | (4%) | (-3%) | (19%) | (1%) |
| Type B | 0 | 30 | 2 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | -1 |
| | N/A | (63%) | (6%) | (0%) | (0%) | (0%) | (0%) | (0%) | (0%) | (0%) | (0%) | (0%) |

Percentage change in given time period in parentheses. $\Delta_{p_n}$ and $\Delta_{p_r}$ represent the change in new purchase and renewal price, respectively, such that $\Delta_{p_n} / \Delta_{p_r} =$ -\$1 / -\$1 means that there is a \$1 decrease in both the new purchase and renewal price in that month. As in the model specification, $n$, $r$ and $g$ represent the number of new purchases, renewals and upgrades, respectively. "Total" stands for the aggregate change in new purchases, renewal and upgrades, while "Type A" and "Type B" further breaks these changes out by consumers type.

From Table 6, we see that there is an immediate response to a \$1 discount off the purchase and renewal prices. In particular, in the month of the discount, there is a large increase in the number of renewals (62 or 117%) and a smaller increase in the number of new purchases (3 new purchases, up from 0). The cumulative change in renewals in the following month is much smaller (33 more renewals relative to the baseline case or 30% increase), indicating that the one-time price change creates intertemporal renewal substitution. That is, consumers who would have renewed in the following month at a higher price, decide to renew today to take advantage of the discount. On the other hand, the cumulative renewal response 12 and 18 months after the discount reveals that, in addition to pulling some renewals forward, the discount also creates new renewals. In fact, 18 months after the price change, there are 26 more renewals relative to the baseline case; however, this increase represents a relatively small percentage change (2% increase) as compared the total number of renewals in 18 months. These findings are in-line with previous literature on dynamic consumer response to temporary price reduction, which documents that short-run price elasticities overestimate response to temporary price shocks (Erdem, Imai and Keane 2003, Hendel and Nevo 2006).

I further explore the particularities of this dynamic response in my context by examining the response among the two consumer types separately. The comparison between Type A and Type B consumer response in rows 2 and 3 further shows that Type A consumers respond to the price change by purchasing more. In fact, Type A consumers are the ones that drive the long-run response to the price change. Since Type A consumers are price sensitive, a price discount causes them to renew when they otherwise would not have. Conversely, Type B consumers drive the intertemporal substitution in renewal activity. These consumers value the product enough to continue renewing the subscription even without the discount, but the discount provides additional incentive to renew today rather than in the following month.

Another interesting feature is that a change in either the new purchase or renewal price, but not both, creates a shift between new purchases and renewals. For instance, in the bottom section, when the firm gives a $1 discount on the renewal price, some Type A consumers shift from purchasing a new subscription license to renewing the existing subscription license instead. On the other hand, when the firm gives a $1 discount on the new purchase price, both consumer types shift from renewing the existing license to purchasing a new subscription. I discuss this shift in more detail in the following section.

# 8 Counterfactual Simulations

## 8.1 Scenario 1: High Types Value Innovation

The previous results suggest that by putting out new versions of the product, the firm appears to cater towards the broader market, who value newness and a low price, to the detriment of its highest value consumers, who are willing to pay more for older versions of the product[55]. This feature may be specific to the software industry, where in many cases, the core functionality of the software product is not improving drastically, but rather, the innovation is around the look or speed of the product function. If the high value consumers do not value this type of innovation, it might be unprofitable for the firm to charge higher prices for the newest functionality; however, this new functionality can attract the low type consumers who otherwise would be priced out.

First, I consider a counterfactual scenario in which the firm offers the newest functionality at a premium. One manner of implementing this type of pricing is to charge the consumer a price to

[55]This outcome is the flip-side to the innovator's dilemma described by Christensen (1997). According to Christensen, when firms cater to the current needs of their most valuable consumers, they often fail to introduce disruptive innovation that caters to untapped markets or consumers' future needs.

upgrade to the newest product version[56]. Thus, in this counterfactual scenario, I also introduce an upgrade price, which enters into the consumer's current period utility from a choice to upgrade as well as his continuation value of future choices.

The left panel of Figure 7 shows that upgrade pricing would lead to a decrease in product upgrades and new subscription re-purchases: the higher the upgrade price the lower the number of upgrades and number of new purchases. The number of renewals would slightly increase, since at higher upgrade prices, consumers prefer to renew an older version of the product without upgrading, and the high value consumers who renew value older versions of the product in any case. Moreover, as shown by the yellow line in Figure 8, the firm's profits with any positive upgrade price would be lower than in the baseline current case with free upgrades.

Next, I examine a counterfactual scenario under which a firm innovates in a way that is appealing to the high value consumer and discuss the effects of introduction of upgrade pricing. In particular, I consider the case when the high type's use value for versions 1 and 3 are reversed; i.e., a Type B consumer's use values for versions 1, 2 and 3 are 3.6699, 3.7906 and 3.9348, respectively. Since the data context is anonymized, I cannot measure how consumers may respond to particular features of different product versions; however, in this counterfactual scenario, I assume that, taking into account all the product feature changes from version 1 to version 3, the high value consumer receives an increasingly positive use utility from each successive version. One way the firm could introduce such a scenario is by discontinuing intra-version security patches. In this case, legacy versions become less appealing to all consumers.

The illustrative model of the firm's pricing decisions in this context (in Appendix Section H) suggests that the firm may want to price the new functionality at a premium. Indeed, Figures 7 and 8 confirm that the firm benefits from pricing the upgrades. As shown in the right panel of Figure 7, when high value consumers value upgrades and upgrades are free, there are fewer renewals, but more upgrades and new purchases, as compared to the current baseline case. From Figure 8, however, we see that when upgrades are free, the firm's profits are lower than in the baseline case. This is because fewer consumers choose the more costly renewal option when a cheaper option that gives them the desired new functionality is available. When the firm charges an upgrade price, the number of upgrades decrease accordingly; however, the firm's profits grow because of the growing upgrade price. At high upgrade prices (15%+ of the more renewal price), consumers begin to shift away even more aggressively from renewals of existing subscription licenses towards purchases of new licenses with the latest and greatest functionality, but a lower price. To the firm, this shift again represents a loss, since the new license price is lower than the renewal price. In fact, as evidenced by the shape of the firm's profits as a function of the upgrade price in Figure 8, at these high upgrade prices, the firm's profits plateau and then begin to drop.

The purpose of evaluating this counterfactual scenario is to illustrate that the firm's ability to

[56] Another manner of implementing this type of pricing would be to flip the new subscription license purchase price and renewal price. In the given data context, the low new purchase price may function partially to attract new consumers. Since I do not estimate the arrival of first purchases, I will not be able to account for the loss of consumers due to a higher new purchase price. As a result, instead of implementing this type premium pricing of new functionality, I introduce an upgrade price.

Figure 7: Consumer Actions as a Function of Upgrade Price

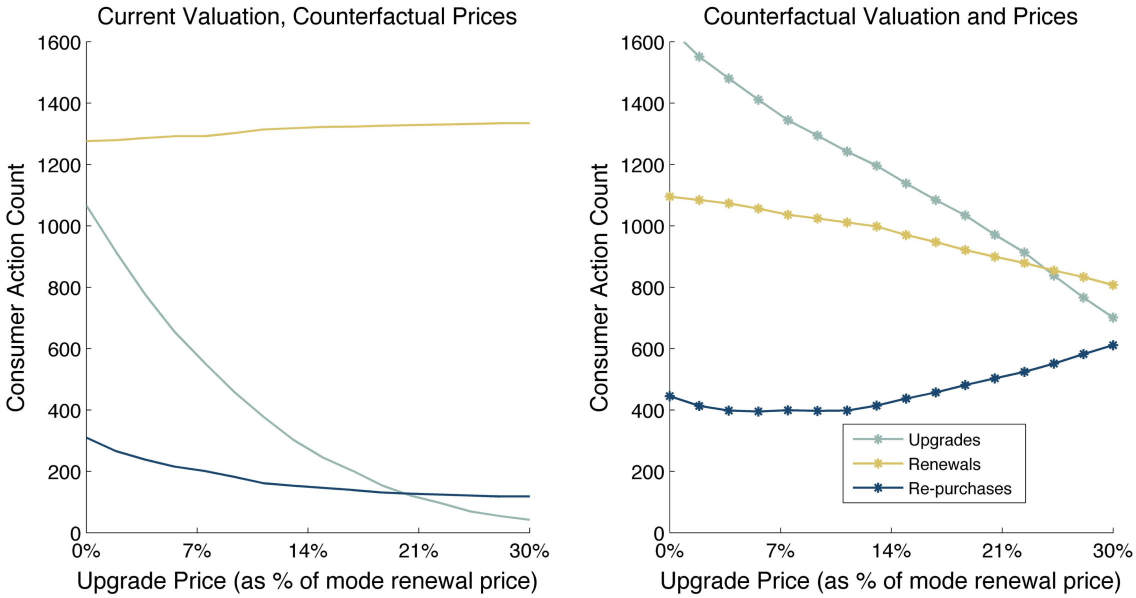


Parameter Values: Heterogeneous Model

Figure 8: Firm Profit as a Function of Upgrade Price

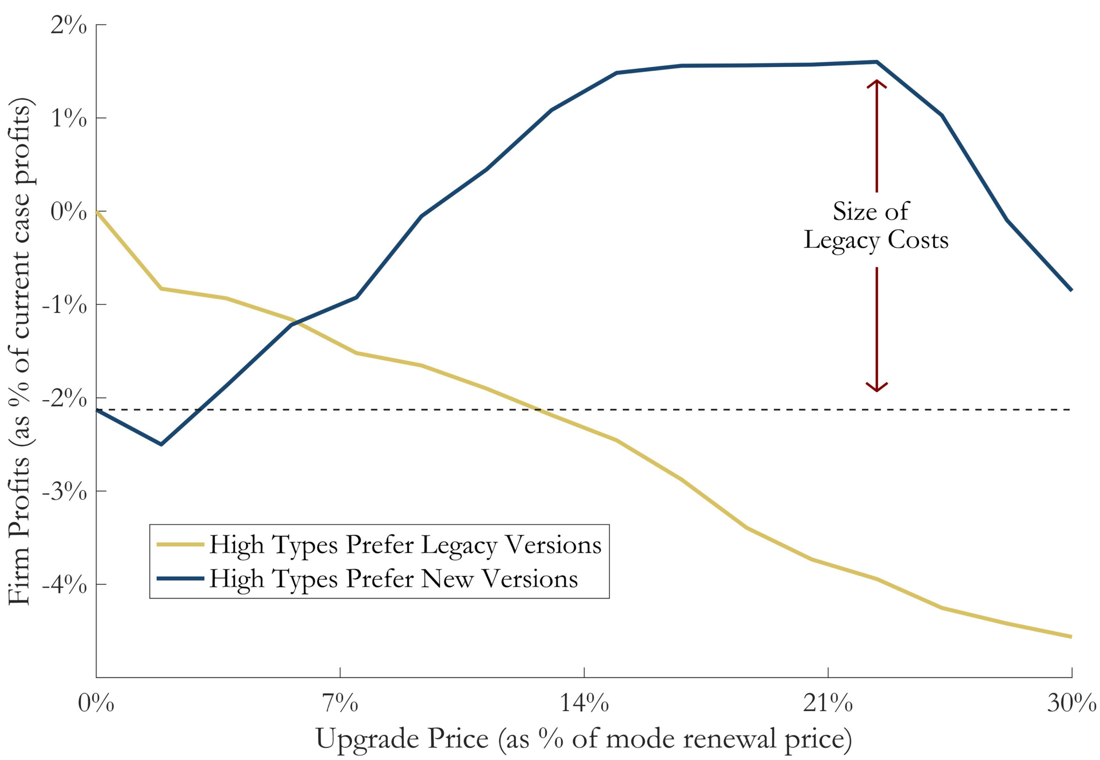


Parameter Values: Heterogeneous Model

price new functionality at a premium depends on the extent that high value consumers value the new functionality over old. In this empirical application, it is optimal for the firm to offer free upgrades because high value consumers do not value the innovation; however, if the firm were to introduce new version innovation in a way that is valued by high value consumers, charging a zero upgrade price becomes sub-optimal. In the following sections, I further investigate the extent to which costly software product maintenance may additionally prevent the software firm from price discriminating on the newest functionality, even when new high value consumers value the firm's innovation.

## 8.2 Scenario 2: Costly Legacy Versions and Innovation (In Progress)

As compared to the typical durable goods firm, the software firm is unique in that it may have costs associated with its software product, even after its sale. This issue is particularly relevant in operating systems and security software, where the firm is expected to continue to issue security patches to the product to manage any vulnerabilities. Moreover, compared to current versions, legacy versions may be more costly to service due to specialized manpower needed to deal with legacy code, for instance. Legacy versions may also be more vulnerable to piracy, which further and disproportionately drives up their costs from the firm's perspective.

In the previous counterfactual scenario, I showed that if the firm were able to introduce innovation that is valued by the high type consumers to some extent, it could increase its profit by charging an upgrade price. In this section, I consider how, even in this case, the size of costs associated with legacy versions may limit the firm's ability to price discriminate on newest functionality. In particular, I consider a scenario in which legacy version costs represent a particular portion of the firm's profit (e.g., 5%)[57]. Given this cost, I evaluate how much the firm has to innovate in order for it to be able to price discriminate on newest functionality. To implement this, I consider different increments of use value from one version to the next for the high type consumer. Given this level of innovation, I then solve for the optimal upgrade price and compare the firm's gains from paid product upgrades to legacy version costs. This counterfactual exercise allows me to further quantify the restrictions on the firm's ability to price discriminate on the newest version in presence of legacy costs.

## 8.3 Scenario 3: Phasing Out Costly Legacy Versions (In Progress)

In this section, I address a major current trend in the software industry: a shift away from maintaining a number of intertemporal product versions. Several different software firms, including Microsoft Office and Adobe Creative Cloud, are shifting towards a one-version software product that receives free upgrades to the latest functionality. Moreover, some firms are forcing upgrade compliance by eliminating intertemporal versions altogether (e.g., Adobe 2013), rolling out automatic upgrades (e.g., Microsoft Windows via Keizer 2016) or discontinuing support for legacy versions (e.g., MATLAB via Mathworks 2016). By taking these actions, firms are reducing the number of

[57] In discussions with the data sponsor firm to confirm the appropriate legacy version costs to use for the calculation.

legacy versions in the market and, consequently, the costs associated with maintaining these versions. Such a policy may adversely impact consumers who value old product versions, however, and have a negative impact on the firm's profits.

In this counterfactual exercise, I evaluate the impact of a shift away from legacy versions in my empirical context. Instead of making upgrades optional, I introduce a scenario in which the firm implements a mandatory upgrade to the newest version as soon as the version is introduced, and the consumer takes this policy into account when making purchase and renewal decisions. As shown in Section 7, in this empirical application, the high type consumers value older versions of the product more than the newest functionality, thus driving legacy version usage. If these consumers are then forced to upgrade to the newest functionality, they receive lower use value from the product than they would have if they were allowed to remain on the older version of the product. As a consequence, after the subscription with mandatory upgrades expires, some of these consumers may choose not to renew the license, even though they would have in the baseline case. The firm suffers losses in revenue from the high types; however, it also eliminates the costs associated with legacy versions. This counterfactual exercise allows me to quantify how high the legacy version maintenance costs would have to be in order for the firm to shift away entirely from maintaining intertemporal versions of the product.

# 9 Conclusions and Future Work

To my knowledge, this paper is the first to empirically study the factors affecting the shift away from intertemporal price discrimination via versioning and intertemporal versions altogether. In my analysis, I use a novel WCAI data grant that tracks consumers' purchases, renewals and upgrades of the sponsoring firm's products. Through descriptive analysis, I document that the consumers in this empirical context frequently choose not to upgrade to the latest version and continue to renew legacy versions of the product. I then specify and estimate a dynamic structural model of consumers choice to show that the usage of legacy versions is driven by the firm's high value, price insensitive consumers. Through a counterfactual simulation, I show that, although version price discrimination is not optimal in the current scenario, in the more typical durable goods scenario where high value consumers value the firm's innovation, it is more profitable for the firm to price discriminate.

In the coming months, I plan to complete the in-progress counterfactual simulations that relate the firm's ability to price discriminate to the costs of maintaining legacy versions. In one counterfactual scenario, I follow-up on the finding that price discrimination is profitable if high value consumers value innovation. I examine how much a firm needs to innovate in order for price discrimination to be profitable if the firm has non-negligible costs of maintaining legacy product versions. The intuition behind this counterfactual is that increased innovation allows the firm to price newer versions at a premium; however, this type of pricing may also cause more consumers to remain on legacy versions, which contributes to the firm's costs. Thus, a firm must ensure that the innovation is above a certain threshold in order to profitably price discriminate on newest functionality. In a second counterfactual scenario, I examine the extent to which legacy costs have to be limiting in order for the firm to shift away entirely from legacy versions towards one continuously maintained version. The intuition behind this counterfactual is that, even though the high value, low price sensitivity consumers value the legacy versions, the costs of servicing these versions can be high enough for the firm to eliminate them altogether. I compute the minimum level of the costs associated with legacy versions that would induce a firm to make this move.

Future work could extend the work presented here in two ways. First, it could supplement the model of individual consumer behavior by considering the consumers' first purchases. The current dataset is comprised of five roughly equal consumer cohorts, and the purchase of the machine on which the software is installed is unobserved. It may be possible to model first purchases using aggregate first purchase data, however, additional information is required from the sponsoring firm to do so. Alternatively, the first instance of free trial or other product purchase could be used to infer the purchase of the machine. Both of these approaches would further inform the extent to which different consumer types choose to first purchase a given version of a product over another. Future work could also address the firm's innovation process and the consumers' uncertainty over the use values of the product versions. In my current model, consumers have perfect foresight of the use values of each of the successive versions. It is likely, however, that firms cannot commit exactly to a particular version quality, especially with a fixed version release schedule. Rather, the quality of a particular version depends on the time and resources the firm spends on innovation as well as time-varying shocks to these costs. With additional data on the firm's innovation process, it would be possible to model this quality uncertainty, which would help inform whether this particular firm is optimally choosing a subscription rather than a perpetual license product offering.

# 10 References

"A Practical Guide to Legacy Application Retirement. Archiving Data with the Informatica Solution for Application Retirement." informatica.com, September 2013. Informatica. Web. 29 May 2016. <https://www.informatica.com/content/dam/informatica-com/global/amer/us/collateral/white-paper/practical-guide-legacy-application-retirement_white-paper_6957.pdf>.

"Adobe's Shift to the Cloud: Is This the Start of a Trend?." Knowledge@Wharton. The Wharton School, University of Pennsylvania, 08 May 2013. Web. 28 May 2016 <http://knowledge.wharton.upenn.edu/article/adobes-shift-to-the-cloud-is-this-the-start-of-a-trend/>.

Adobe. "Adobe Accelerates Shift to the Cloud". adobe.com., 6 May 2013. Web. 28 May 2016. <http://www.adobe.com/aboutadobe/pressroom/pressreleases/201305/050613 AdobeAcceleratesShifttotheCloud.html>.

Albuquerque, Paulo, and Yulia Nevskaya. "The Impact of Innovation and Social Interactions on Product Usage." Working Paper, [available at http://server1. tepper. cmu. edu/seminars/docs/ Product_Usage_Paper. pdf] (2015).

Bhargava, Hemant K., and Vidyanand Choudhary. "Research Note-When Is Versioning Optimal for Information Goods?." Management Science 54.5 (2008): 1029-1035.

Borkovsky, Ron N., "The Timing of Version Releases: A Dynamic Duopoly Model." Working Paper (2015).

Bronnenberg, Bart J. "The provision of convenience and variety by the market." The RAND Journal of Economics 46.3 (2015): 480-498.

Bulow, Jeremy. "An economic theory of planned obsolescence." The Quarterly Journal of Economics (1986): 729-750.

Clayton, Christensen. "The innovator's dilemma: When new technologies cause great firms to fail." Harvard Business School Print (1997).

Desai, Preyas S. "Quality segmentation in spatial markets: When does cannibalization affect product line design?." Marketing Science 20.3 (2001): 265-283.

Desai, Preyas, et al. "Product differentiation and commonality in design: Balancing revenue and cost drivers." Management Science 47.1 (2001): 37-51.

Draganska, Michaela, and Dipak C. Jain. "Consumer preferences and product-line pricing strategies: An empirical analysis." Marketing science 25.2 (2006): 164-174.

Deneckere, Raymond J., and R. Preston McAfee. "Damaged goods." Journal of Economics & Management Strategy 5.2 (1996): 149-174.

Dhebar, Anirudh. "Durable-goods monopolists, rational consumers, and improving products." Marketing Science 13.1 (1994): 100-120.

Ellison, Glenn, and Drew Fudenberg. "The neo-Luddite's lament: Excessive upgrades in the software industry." The RAND Journal of Economics (2000): 253-272.

Erdem, Tülin, Susumu Imai, and Michael P. Keane. "Brand and quantity choice dynamics under price uncertainty." Quantitative Marketing and Economics 1.1 (2003): 5-64.

Fishman, Arthur, and Rafael Rob. "Product innovation by a durable-good monopoly." The RAND journal of Economics (2000): 237-252.

Fudenberg, Drew, and Jean Tirole. "Upgrades, tradeins, and buybacks." The RAND Journal of Economics (1998): 235-258.

Goettler, Ronald L., and Karen Clay. "Tariff choice with consumer learning and switching costs." Journal of Marketing research 48.4 (2011): 633-652.

Goettler, Ronald L., and Brett R. Gordon. "Does AMD spur Intel to innovate more?." Journal of Political Economy 119.6 (2011): 1141-1200.

Gordon, Brett R. “A dynamic model of consumer replacement cycles in the PC processor industry.” Marketing Science 28.5 (2009): 846-867.

Gowrisankaran, Gautam, and Marc Rysman. “Dynamics of Consumer Demand for New Durable Goods.” Journal of Political Economy 120.6 (2012): 1173-1219.

Hauser, John, Gerard J. Tellis, and Abbie Griffin. “Research on innovation: A review and agenda for marketing science.” Marketing science 25.6 (2006): 687-717.

Hartmann, Wesley R., and Harikesh S. Nair. “Retail competition and the dynamics of demand for tied goods.” Marketing Science 29.2 (2010): 366-386.

Hendel, Igal, and Aviv Nevo. “Measuring the implications of sales and consumer inventory behavior.” Econometrica 74.6 (2006): 1637-1673.

Hoffman, Chris. “No, Windows 10 Won’t Require a Subscription: Here’s How Microsoft Plans on Making Money Instead.” howtogeek.com., 6 July 2015. Web. 27 February 2016 <http://www.howtogeek.com/220957/no-windows-10-won%E2%80%99t-require-a-subscription-heres-how-microsoft-plans-on-making-money-instead/>.

Keizer, Gregg. “Microsoft makes final, aggressive Windows 10 upgrade push.” computerworld.com., 16 May 2016. Web. 16 June 2016 <http://www.computerworld.com/article/3070460/windows-pcs/microsoft-makes-final-aggressive-windows-10-upgrade-push.html>.

Levinthal, Daniel A., and Devavrat Purohit. “Durable goods and product obsolescence.” Marketing Science 8.1 (1989): 35-56.

Li, Xing. “Product Offerings and Product Line Length Dynamics.” Available at SSRN 2462577 (2014).

Microsoft. Microsoft News Center. “Microsoft Launches Office 365 Globally”. news.microsoft.com. N.p., 28 June 2011. Web. 28 May 2016. <https://news.microsoft.com/2011/06/28/microsoft-launches-office-365-globally/>.

Microsoft. Microsoft News Center. “Windows 10 Available as a Free Upgrade on July 29”. news.microsoft.com. N.p., 1 July 2015. Web. 28 May 2016. <https://news.microsoft.com /2015/06/01/windows-10-available-as-a-free-upgrade-on-july-29/>.

Nair, Harikesh. “Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games.” Quantitative Marketing and Economics 5.3 (2007): 239-292.

Netessine, Serguei, and Terry A. Taylor. “Product line design and production technology.” Marketing Science 26.1 (2007): 101-117.

Narayanan, Sridhar, Pradeep K. Chintagunta, and Eugenio J. Miravete. "The role of self selection, usage uncertainty and learning in the demand for local telephone service." Quantitative Marketing and Economics 5.1 (2007): 1-34.

Peres, Renana, Eitan Muller, and Vijay Mahajan. "Innovation diffusion and new product growth models: A critical review and research directions." International Journal of Research in Marketing 27.2 (2010): 91-106.

Rust, John. "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher." Econometrica: Journal of the Econometric Society (1987): 999-1033.

"Update to New MATLAB releases and more." mathworks.com., Web. 16 June 2016 <https://www.mathworks.com/services/maintenance/>.

Varian, Hal R. "Versioning Information Goods." University of California, Berkeley (1997).

Viard, V. Brian. "Information goods upgrades: theory and evidence." The BE Journal of Theoretical Economics 7.1 (2007).

Villas-Boas, J. Miguel. "Communication strategies and product line design." Marketing Science 23.3 (2004): 304-316.

Zhang, Jie, and Abraham Seidmann. "Perpetual versus subscription licensing under quality uncertainty and network externality effects." Journal of management information systems 27.1 (2010): 39-68.

# Appendix

## A Price Panels

The price panels presented here are obtained using the 2-step method described in Section 3.

Figure 9: New Purchase Prices by Product and Subscription Length

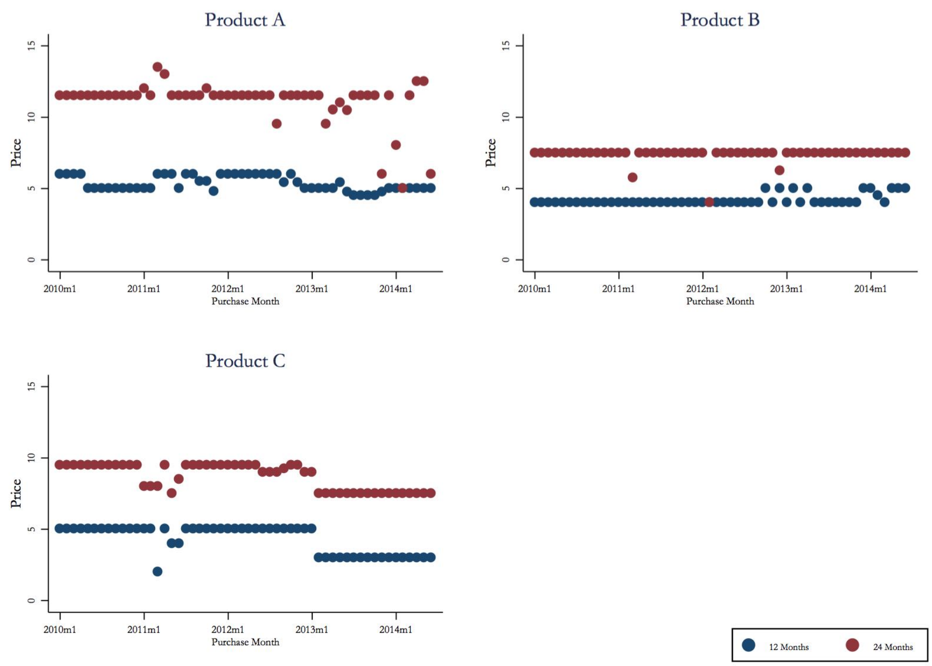


Figure 10: Renewal Prices by Product and Subscription Length

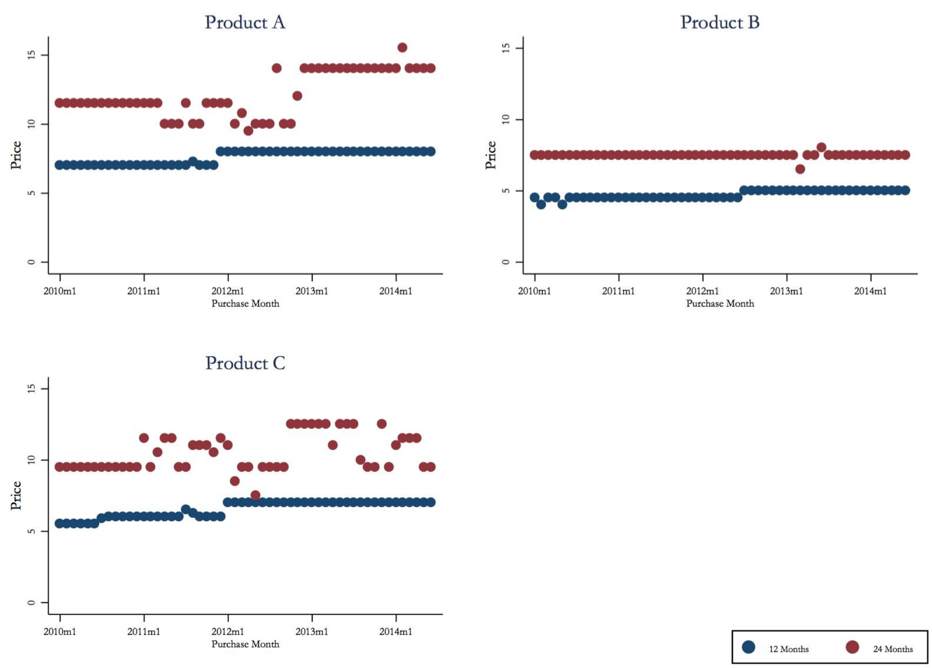

# B Descriptive Evidence, Products B and C

Figure 11: Paid New Purchases and Renewals by Version (Product B)

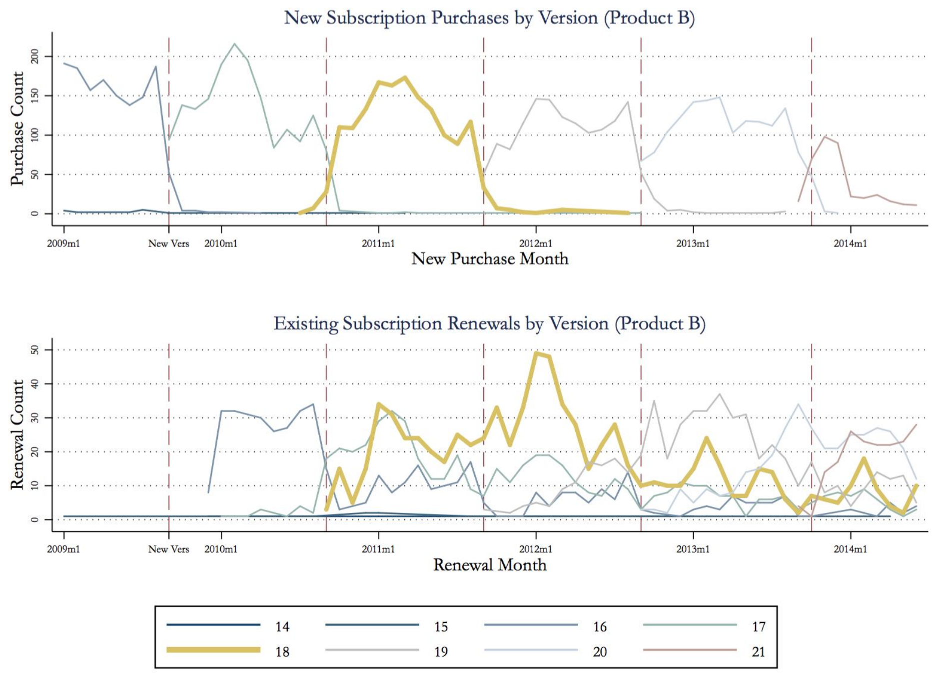


Note: graph includes paid purchases and renewals only, excludes trials.

Figure 12: Paid New Purchases and Renewals by Version (Product C)

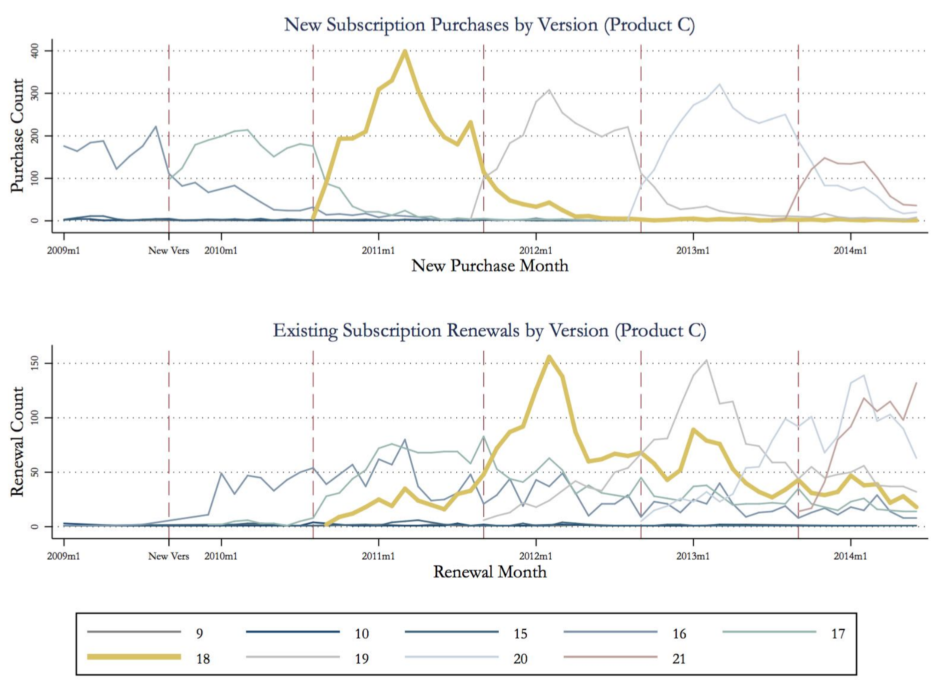


Note: graph includes paid purchases and renewals only, excludes trials.

Figure 13: Upgrades to a Given Product Version (Product B)

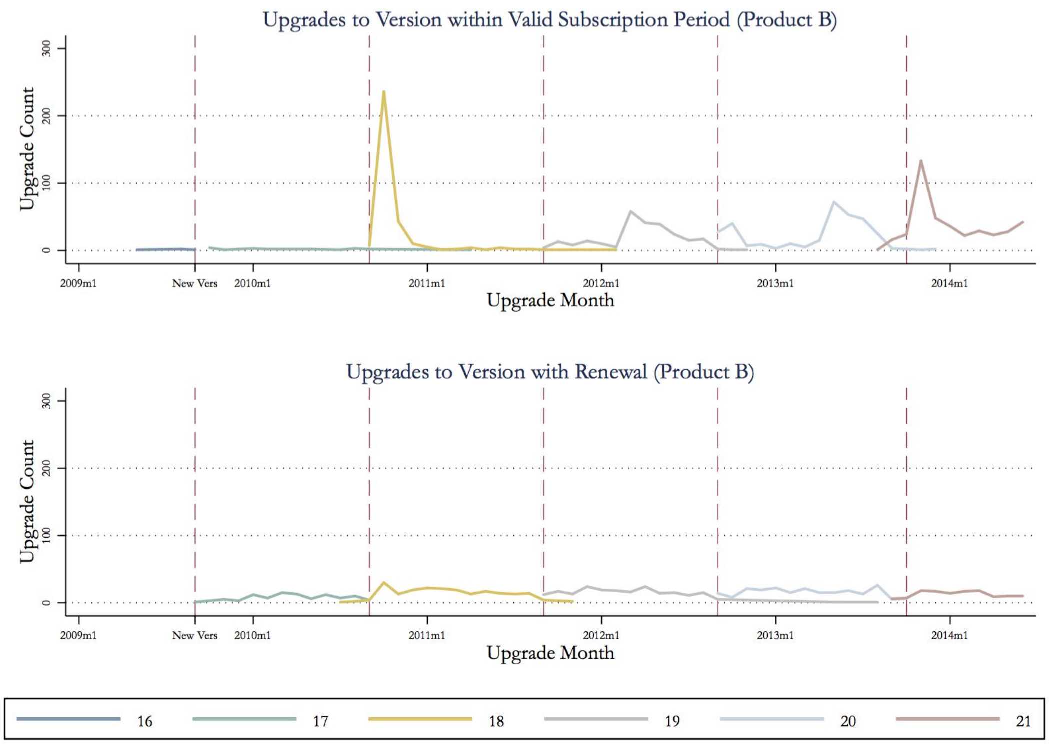


Figure 14: Upgrades to a Given Product Version (Product C)

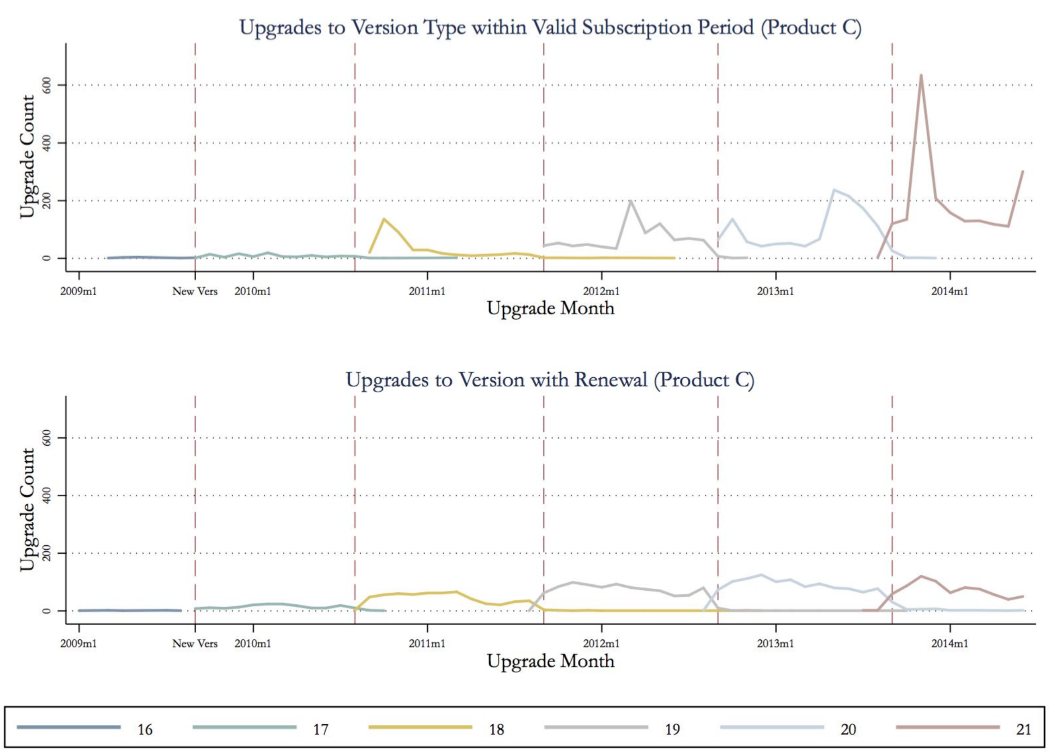

Table 7: Proportion of Eligible Consumers Upgrading (Product B)

| | Newest Available Version | | | | |
|---|---|---|---|---|---|
| Current Version | 17 | 18 | 19 | 20 | 21 |
| 16 | 0.06 | 0.31 | 0.05 | 0.10 | 0.00 |
| 17 | . | 0.20 | 0.09 | 0.12 | 0.06 |
| 18 | . | . | 0.19 | 0.19 | 0.06 |
| 19 | . | . | . | 0.22 | 0.15 |
| 20 | . | . | . | . | 0.25 |
| All | 0.06 | 0.22 | 0.17 | 0.21 | 0.21 |

Table 8: Proportion of Eligible Consumers Upgrading (Product C)

| | Newest Available Version | | | | |
|---|---|---|---|---|---|
| Current Version | 17 | 18 | 19 | 20 | 21 |
| 16 | 0.07 | 0.18 | 0.12 | 0.10 | 0.04 |
| 17 | . | 0.13 | 0.15 | 0.09 | 0.05 |
| 18 | . | . | 0.25 | 0.19 | 0.12 |
| 19 | . | . | . | 0.34 | 0.19 |
| 20 | . | . | . | . | 0.42 |
| All | 0.07 | 0.15 | 0.22 | 0.27 | 0.33 |

Figure 15: Difference in Time to Expiration at Upgrade in / out of Firm Communication Months

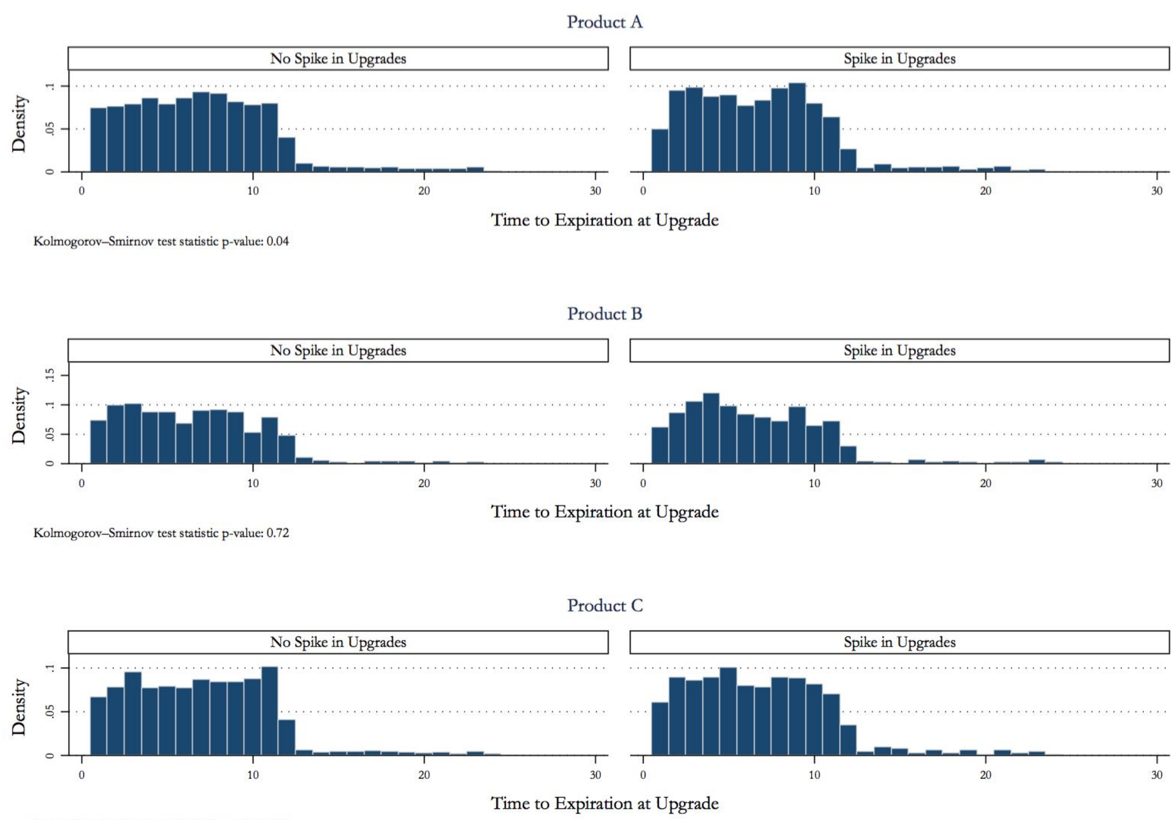

# C Choice-Specific Value Functions in All States

Using the states described in Table 4 as indices, I write the consumer's choice-specific value functions in all possible states of the world:

### Never Purchased

1. **State** $m = 0$**,** $e = 0$**; Choices** $a = (n, r, g) \in \{(1,0,0), (0,0,0)\}$:
   The choices available are purchase of new subscription license ($n$) and no action.

   - $s > 1$:

$$v_{1,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max\big[ v_{2,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \; v_{2,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1} \big] \Big], \text{if } a = (1,0,0); \tag{17}$$

$$v_{1,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max\big[ v_{1,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \; v_{1,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1} \big] \Big], \text{if } a = (0,0,0).$$

   - $s = 1$:

$$v_{1,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max\big[ v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \; v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \; v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1} \big] \Big], \text{ if } a = (1,0,0); \tag{18}$$

$$v_{1,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max\big[ v_{1,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \; v_{1,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1} \big] \Big], \text{if } a = (0,0,0).$$

### Valid Subscription

2. **State** $m > 1$**,** $e = \bar{e}$**; Choices** $a = (n, r, g) \in \{(1,0,0), (0,0,0)\}$:
   The choices available are purchase of new subscription license ($n$) and no action.

   - $m > 2$, $s > 1$:

$$v_{2,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max\big[ v_{2,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \; v_{2,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1} \big] \Big], \; \forall a \in \{(1,0,0), (0,0,0)\}. \tag{19}$$

   - $m > 2$, $s = 1$:

$$v_{2,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon \Big[ \max\big[ v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1},$$

$$
\begin{aligned}
& v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \qquad (20)\\
& v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\Big]\Big], \ \forall a \in \{(1,0,0),(0,0,0)\}.
\end{aligned}
$$

- $m = 2$, $s > 1$:

$$
\begin{aligned}
v_{2,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{2,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1},\\
& v_{2,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{if } a = (1,0,0);\\
v_{2,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{4,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \qquad (21)\\
& v_{4,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1},\\
& v_{4,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (0,0,0).
\end{aligned}
$$

- $m = 2$, $s = 1$:

$$
\begin{aligned}
v_{2,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1},\\
& v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1},\\
& v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (1,0,0);\\
v_{2,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{5,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \qquad (22)\\
& v_{5,(0,1,1)}(x_{t+1}) + \varepsilon_{(0,1,1),t+1},\\
& v_{5,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1},\\
& v_{5,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1},\\
& v_{5,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (0,0,0).
\end{aligned}
$$

3. **State** $m > 1$, $e < \bar{e}$**; Choices** $a = (n, r, g) \in \{(1,0,0),(0,0,1),(0,0,0)\}$:
   The choices available are purchase of new subscription license ($n$), product upgrade ($g$) and no action.

   - $m > 2$, $s > 1$:

$$
\begin{aligned}
v_{3,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{2,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1},\\
& v_{2,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \ \forall a \in \{(1,0,0),(0,0,1)\};\\
v_{3,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \qquad (23)\\
& v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1},
\end{aligned}
$$

$$v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\Big]\Big], \text{ if } a = (0,0,0).$$

- $m > 2$, $s = 1$:

$$\begin{aligned} v_{3,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\ & v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \\ & v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \ \forall a \in \{(1,0,0),(0,0,1),(0,0,0)\}. \end{aligned} \tag{24}$$

- $m = 2$, $s > 1$:

$$\begin{aligned} v_{3,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{2,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\ & v_{2,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{if } a = (1,0,0); \\ v_{3,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{4,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\ & v_{4,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1}, \\ & v_{4,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (0,0,1); \\ v_{3,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{5,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\ & v_{5,(0,1,1)}(x_{t+1}) + \varepsilon_{(0,1,1),t+1}, \\ & v_{5,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1}, \\ & v_{5,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \\ & v_{5,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (0,0,0). \end{aligned} \tag{25}$$

- $m = 2$, $s = 1$:

$$\begin{aligned} v_{3,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\ & v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \\ & v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (1,0,0); \\ v_{3,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{5,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\ & v_{5,(0,1,1)}(x_{t+1}) + \varepsilon_{(0,1,1),t+1}, \\ & v_{5,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1}, \\ & v_{5,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \end{aligned} \tag{26}$$

$$v_{5,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\Big]\Big], \ \forall a \in \{(0,0,1),(0,0,0)\}.$$

4. **State** $m = 1$, $e = \bar{e}$; **Choices** $a = (n, r, g) \in \{(1,0,0),(0,1,0),(0,0,0)\}$:
   The choices available are purchase of new subscription license ($n$), renewal of existing subscription license ($r$) and no action.

   - $s > 1$:

$$\begin{aligned}
v_{4,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{2,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{2,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \ \forall a \in \{(1,0,0),(0,1,0)\}; \\
v_{4,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{6,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{6,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1}, \\
& v_{6,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \ \text{if } a = (0,0,0).
\end{aligned} \tag{27}$$

   - $s = 1$:

$$\begin{aligned}
v_{4,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \\
& v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \ \forall a \in \{(1,0,0),(0,1,0)\}; \\
v_{4,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{7,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{7,(0,1,1)}(x_{t+1}) + \varepsilon_{(0,1,1),t+1}, \\
& v_{7,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1}, \\
& v_{7,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \ \text{if } a = (0,0,0).
\end{aligned} \tag{28}$$

5. **State** $m = 1$, $e < \bar{e}$; **Choices** $a = (n, r, g) \in \{(1,0,0),(0,1,0),(0,0,1),(0,1,1),(0,0,0)\}$:
   The choices available are purchase of new subscription license ($n$), renewal of existing subscription license ($r$), product upgrade ($g$), renewal of existing subscription license and product upgrade ($r$, $g$), and no action.

   - $s > 1$:

$$\begin{aligned}
v_{5,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{2,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{2,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \forall a \in \{(1,0,0),(0,1,1)\}; \\
v_{5,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1},
\end{aligned}$$

$$v_{3,(0,0,1)}(x_{t+1})+\varepsilon_{(0,0,1),t+1},$$

$$v_{3,(0,0,0)}(x_{t+1})+\varepsilon_{(0,0,0),t+1}\Big]\Big], \text{ if } a=(0,1,0);$$

$$v_{5,a}(x_t)=u_a(x_t)+\delta\mathbb{E}_\varepsilon\Big[\max\big[v_{6,(1,0,0)}(x_{t+1})+\varepsilon_{(1,0,0),t+1}, \tag{29}$$

$$v_{6,(0,1,0)}(x_{t+1})+\varepsilon_{(0,1,0),t+1},$$

$$v_{6,(0,0,0)}(x_{t+1})+\varepsilon_{(0,0,0),t+1}\Big]\Big], \text{ if } a=(0,0,1);$$

$$v_{5,a}(x_t)=u_a(x_t)+\delta\mathbb{E}_\varepsilon\Big[\max\big[v_{7,(1,0,0)}(x_{t+1})+\varepsilon_{(1,0,0),t+1},$$

$$v_{7,(0,1,1)}(x_{t+1})+\varepsilon_{(0,1,1),t+1},$$

$$v_{7,(0,1,0)}(x_{t+1})+\varepsilon_{(0,1,0),t+1},$$

$$v_{7,(0,0,0)}(x_{t+1})+\varepsilon_{(0,0,0),t+1}\Big]\Big], \text{ if } a=(0,0,0).$$

- $s=1$:

$$v_{5,a}(x_t)=u_a(x_t)+\delta\mathbb{E}_\varepsilon\Big[\max\big[v_{3,(1,0,0)}(x_{t+1})+\varepsilon_{(1,0,0),t+1},$$

$$v_{3,(0,0,1)}(x_{t+1})+\varepsilon_{(0,0,1),t+1},$$

$$v_{3,(0,0,0)}(x_{t+1})+\varepsilon_{(0,0,0),t+1}\Big]\Big], \forall a\in\{(1,0,0),(0,1,1),(0,1,0)\};$$

$$v_{5,a}(x_t)=u_a(x_t)+\delta\mathbb{E}_\varepsilon\Big[\max\big[v_{7,(1,0,0)}(x_{t+1})+\varepsilon_{(1,0,0),t+1}, \tag{30}$$

$$v_{7,(0,1,1)}(x_{t+1})+\varepsilon_{(0,1,1),t+1},$$

$$v_{7,(0,1,0)}(x_{t+1})+\varepsilon_{(0,1,0),t+1},$$

$$v_{7,(0,0,0)}(x_{t+1})+\varepsilon_{(0,0,0),t+1}\Big]\Big], \forall a\in\{(0,0,1),(0,0,0)\}.$$

## Expired Subscription

6. **State** $m=0$**,** $e=\overline{e}$**; Choices** $a=(n,r,g)\in\{(1,0,0),(0,1,0),(0,0,0)\}$:
   The choices available are purchase of new subscription license ($n$), renewal of existing subscription license ($r$) and no action.

   - $s>1$:

$$v_{6,a}(x_t)=u_a(x_t)+\delta\mathbb{E}_\varepsilon\Big[\max\big[v_{2,(1,0,0)}(x_{t+1})+\varepsilon_{(1,0,0),t+1},$$

$$v_{2,(0,0,0)}(x_{t+1})+\varepsilon_{(0,0,0),t+1}\Big]\Big], \forall a\in\{(1,0,0),(0,1,0)\};$$

$$v_{6,a}(x_t)=u_a(x_t)+\delta\mathbb{E}_\varepsilon\Big[\max\big[v_{6,(1,0,0)}(x_{t+1})+\varepsilon_{(1,0,0),t+1}, \tag{31}$$

$$v_{6,(0,1,0)}(x_{t+1})+\varepsilon_{(0,1,0),t+1},$$

$$v_{6,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\Big]\Big], \text{ if } a = (0,0,0).$$

- $s = 1$:

$$\begin{aligned}
v_{6,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \\
& v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \forall a \in \{(1,0,0),(0,1,0)\}; \\
v_{6,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{7,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{7,(0,1,1)}(x_{t+1}) + \varepsilon_{(0,1,1),t+1}, \\
& v_{7,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1}, \\
& v_{7,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (0,0,0).
\end{aligned} \tag{32}$$

7. **State** $m = 0$, $e < \bar{e}$; **Choices** $a = (n, r, g) \in \{(1,0,0),(0,1,0),(0,1,1),(0,0,0)\}$:
   The choices available are purchase of new subscription license ($n$), renewal of existing subscription license ($r$), renewal of existing subscription license and product upgrade ($r$, $g$) and no action.

- $s > 1$:

$$\begin{aligned}
v_{7,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{2,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{2,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \forall a \in \{(1,0,0),(0,1,1)\}; \\
v_{7,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1}, \\
& v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (0,1,0); \\
v_{7,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{7,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{7,(0,1,1)}(x_{t+1}) + \varepsilon_{(0,1,1),t+1}, \\
& v_{7,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1}, \\
& v_{7,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (0,0,0).
\end{aligned} \tag{33}$$

- $s = 1$:

$$\begin{aligned}
v_{7,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[ \max\big[ & v_{3,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \\
& v_{3,(0,0,1)}(x_{t+1}) + \varepsilon_{(0,0,1),t+1},
\end{aligned}$$

$$v_{3,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\Big]\Big], \forall a \in \{(1,0,0),(0,1,1),(0,1,0)\};$$

$$v_{7,a}(x_t) = u_a(x_t) + \delta \mathbb{E}_\varepsilon\Big[\max\big[v_{7,(1,0,0)}(x_{t+1}) + \varepsilon_{(1,0,0),t+1}, \tag{34}$$
$$v_{7,(0,1,1)}(x_{t+1}) + \varepsilon_{(0,1,1),t+1},$$
$$v_{7,(0,1,0)}(x_{t+1}) + \varepsilon_{(0,1,0),t+1},$$
$$v_{7,(0,0,0)}(x_{t+1}) + \varepsilon_{(0,0,0),t+1}\big]\Big], \text{ if } a = (0,0,0).$$

# D Estimation Product Selection

- TIMEFRAME
  i. Customers who purchase the product
     Product A: **14,081**, Product B: **7,538**, Product C: **19,436**
  ii. Customers who purchase the product starting January 2010 and later (price availability)
     Product A: **10,885**, Product B: **5,581**, Product C: **15,517**
- PURCHASES
  i. Customers who purchase newest version when buying for the first time
     Product A: **9,808**, Product B: **5,280**, Product C: **13,128**
  ii. Customers who purchase a 12 month subscription
     Product A: **8,519**, Product B: **4,738**, Product C: **3,666**
  iii. Customers who do not purchase another product while covered by main product
     Product A: **8,345**, Product B: **4,556**, Product C: **3,551**
  iv. Customers whose purchase results in time-to-expiration consistent with model
     Product A: **8,204**, Product B: **4,491**, Product C: **3,487**
  v. Customers who do not have purchases more than a month before expiration
     Product A: **8,199**, Product B **4,490**, Product C: **3,486**
- UPGRADES
  i. Customers who do not upgrade to newest version before it comes out
     Product A: **8,199**, Product B: **4,480**, Product C: **3,458**
  ii. Customers who upgrade to newest version of product available
     Product A: **8,170**, Product B: **4,476**, Product C: **3,451**

After applying these criteria, arrive at the following set of customers, purchases and upgrades

- % of customers retained of the total purchasing on or after January 2010
  Product A: **75%**, Product B: **80%**, Product C: **22%**
- Number of purchases
  Product A: **13,951**, Product B: **6,225**, Product C: **5,219**

- Number of upgrades
  Product A: **4,582**, Product B: **1,077**, Product C: **1,072**
- % of re-purchases (either purchase of new license or renewal of license)
  Product A: **41%**, Product B: **28%**, Product C: **34%**
- % of re-purchases that are new license purchases
  Product A: **15%**, Product B: **18%**, Product C: **16%**

# E Homogeneous Model Simulations

I carry out a simulation exercise to illustrate that the computation and estimation procedure is well-suited to recover the parameters of interest.

Taking a set of parameters $\Theta = (\theta_{1,1}, \theta_{1,2}, \theta_{1,3}, \theta_2, \theta_3, \theta_4, \theta_5)$ as given, I first simulate the dataset of choices and states for each consumer over time. In particular, I focus on simulating the choices and states after the consumer's first purchase. The simulation is as follows. For each time period, I first update the individual states (time remaining on subscription $m$ and the version of the product last used $e$). Consumers take the aggregate states (newest available version $\bar{e}$, time to next version $s$ and prices $p$) as given in the market. Then, based on the individual states and the aggregate states in the current period, I determine the choices available to the consumer and the value of making each of these choices. I compute the probability of each available choice using the closed-form expression for the probability of a particular action given by the extreme value error assumption (Equation 14). I then compare this probability to a random number $\xi$ from the uniform distribution. For instance, per Table 4, in state 3, there are three possible choices: purchase a new subscription license, upgrade the product to the newest version or do nothing. In this case, I determine the choices in the following manner:

$$a = (n, r, g) = \begin{cases} (1,0,0) & \text{if} \quad \xi < \Pr\{a = (1,0,0)\} \\ (0,0,1) & \text{if} \quad \Pr\{a = (1,0,0)\} \leq \xi < \Pr\{a = (1,0,0)\} + \Pr\{a = (0,0,1)\} \\ (0,0,0) & \text{if} \quad \xi \geq \Pr\{a = (1,0,0)\} + \Pr\{a = (0,0,1)\} \end{cases} \tag{35}$$

Given this choice, I then again update the following period's individual states, and so on.

I simulate the purchase, renewal and upgrade decisions for all of the consumers, conditional on their first purchase. Using different random number generation seeds, I carry out 100 simulations of the data[58]. For each of the simulated datasets, I then run the estimation routine to recover the parameters. The computation and estimation routine performs quite well, as shown in Table 9 and Figure 16, with the estimated parameters not significantly different from the true parameters.

[58]The number of simulations is limited by the somewhat costly computation and estimation routines.

Table 9: Simulation Results around the True Parameters

| | Simulation Parameters | Estimate Mean | Estimate SE | P-Value True = Mean |
|---|---|---|---|---|
| $\theta_{1,1}$ | 0.5000 | 0.5000 | 0.0256 | 0.3989 |
| $\theta_{1,2}$ | 1.0000 | 1.0011 | 0.0249 | 0.3986 |
| $\theta_{1,3}$ | 1.5000 | 1.5024 | 0.0270 | 0.3973 |
| $\theta_2$ | -1.0000 | -1.0013 | 0.0201 | 0.3981 |
| $\theta_3$ | -2.0000 | -2.0027 | 0.0420 | 0.3981 |
| $\theta_4$ | -4.5000 | -4.5152 | 0.0897 | 0.3932 |
| $\theta_5$ | 1.0000 | 1.0020 | 0.0485 | 0.3986 |

Figure 16: Histogram of Parameters Obtained from 100 Simulated Datasets

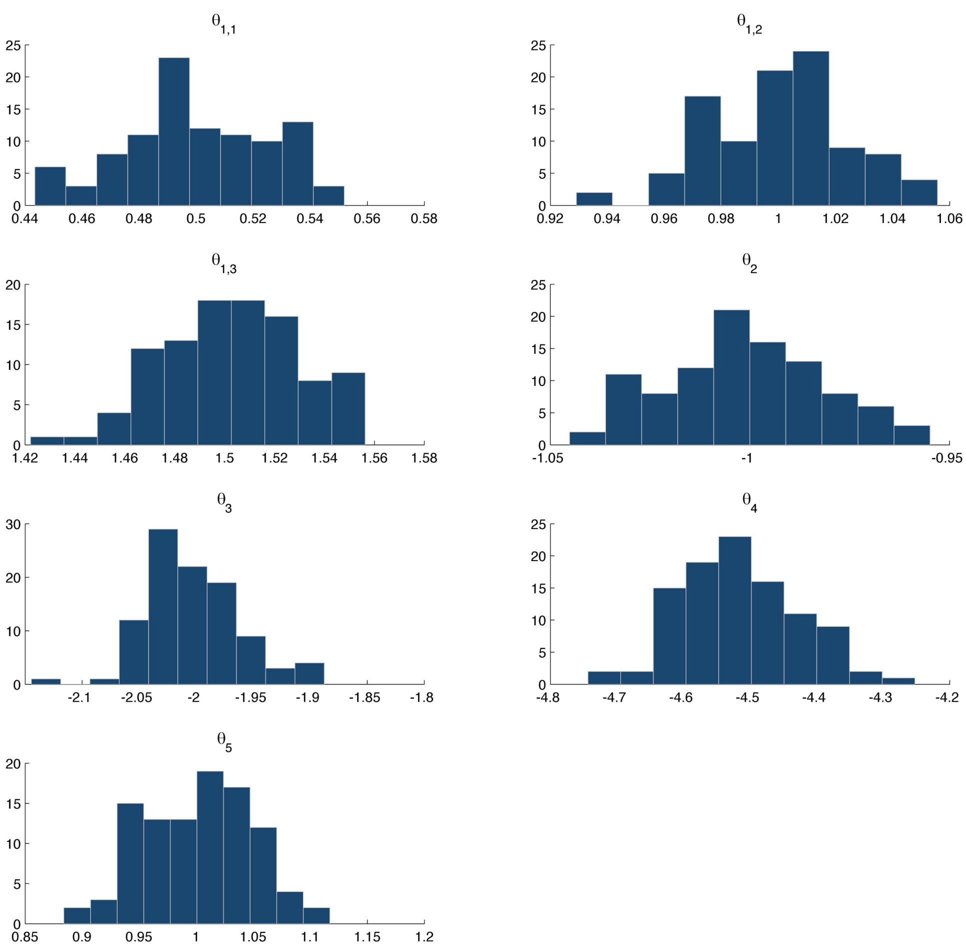

# F Heterogeneous Model Fit

I use the structure of the model and the estimates to simulate the consumers' purchase, renewal and upgrade paths, conditional on the arrival of their first purchases. The aggregate numbers of the different possible actions simulated are as follows:

$$(n,r,g)=\begin{cases}(1,0,0):\text{New Subscription License Purchase}-5,199;\\(0,1,1):\text{Existing Subscription License Renewal and Product Upgrade}-47;\\(0,1,0):\text{Existing Subscription License Renewal}-1,236;\\(0,0,1):\text{Product Upgrade}-1,252;\\(0,0,0):\text{No Action}-69,494.\end{cases}\tag{36}$$

Comparing these aggregate numbers of actions to those found in the data (Equation 16), I note several differences. First, my model understates the number of renewals and upgrades coinciding in the same month. In my model, the utility specification for the renewal and upgrade choice combines the renewal and upgrade elements additively. As a result, in my model, there is no additional benefit to the consumer of taking the renewal and upgrade actions in the same month. One can imagine, however, that, in reality, there is some additional benefit to this coordination; however, given that the difference between the aggregate simulated renewal and upgrade action and the aggregate renewal and upgrade action observed in the data is not very large, I conclude that this additional benefit is fairly small. Conditional on this lack of coordination, the model fits the data quite well overall, as evidenced by Figures 17 and 18. As a result, I proceed with the current model specification and estimates.

Figure 17: Heterogeneous Model Fit: All Actions

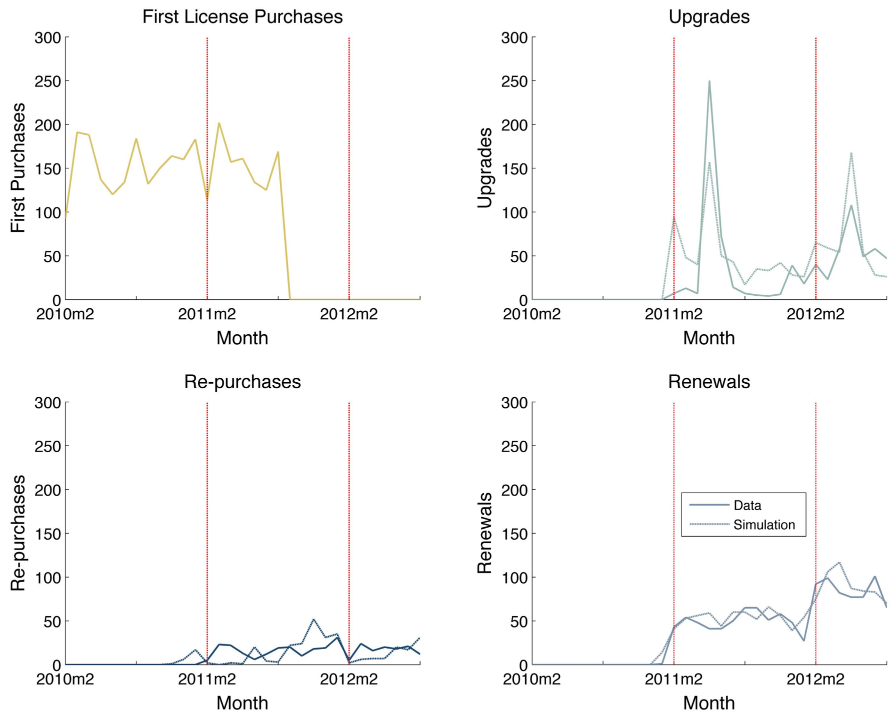


Parameter Values: Heterogeneous Model

Figure 18: Heterogeneous Model Fit: Repurchases by Version

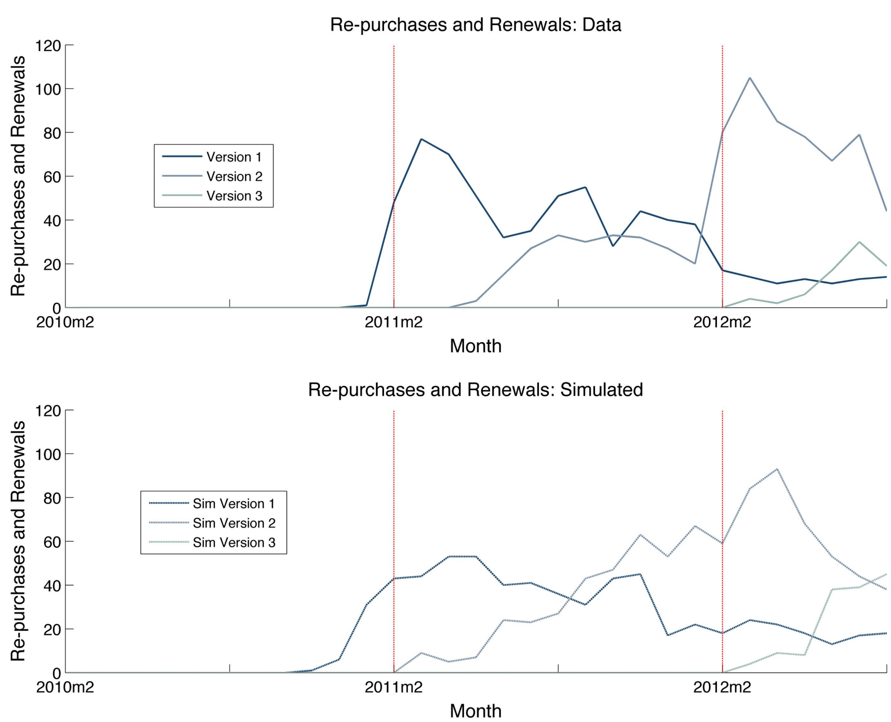


Parameter Values: Heterogeneous Model

# G Value Functions at Estimated Parameters, Select States

Figure 19: Type A Choice Specific Value Functions
at $\bar{e} = 3$, $(p_n, p_r) = (6, 8)$

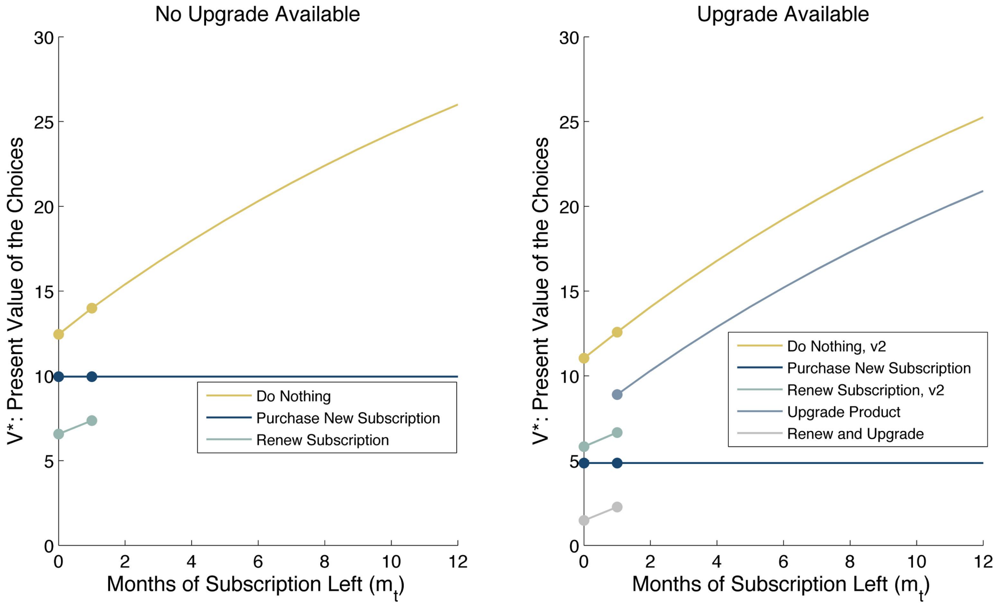


Parameter Values: Heterogeneous Model

Figure 20: Type B Choice Specific Value Functions
at $\bar{e} = 3$, $(p_n, p_r) = (6, 8)$

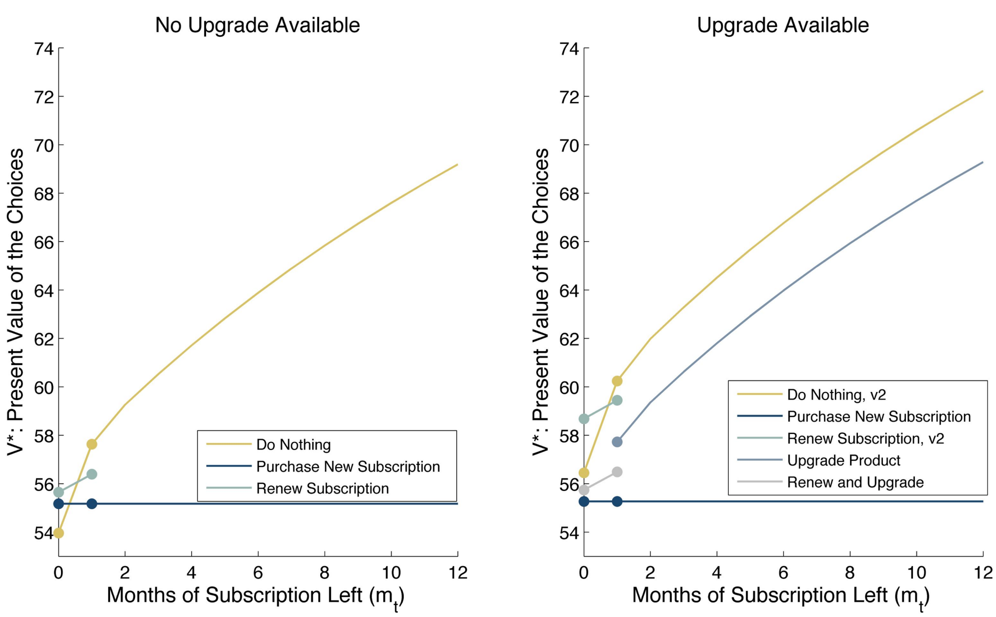


Parameter Values: Heterogeneous Model

Figure 21: Type A Choice Specific Value Functions

at $\bar{e} = 2$, $(p_n, p_r) = (6, 8)$, No Upgrade Available

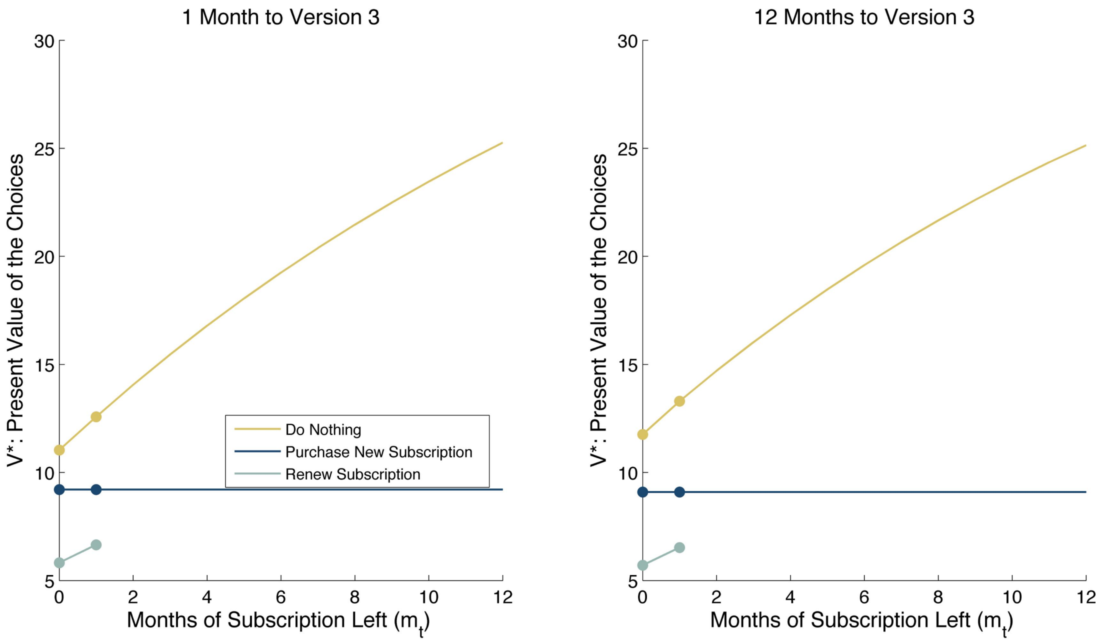


Parameter Values: Heterogeneous Model

Figure 22: Type B Choice Specific Value Functions

at $\bar{e} = 2$, $(p_n, p_r) = (6, 8)$, No Upgrade Available

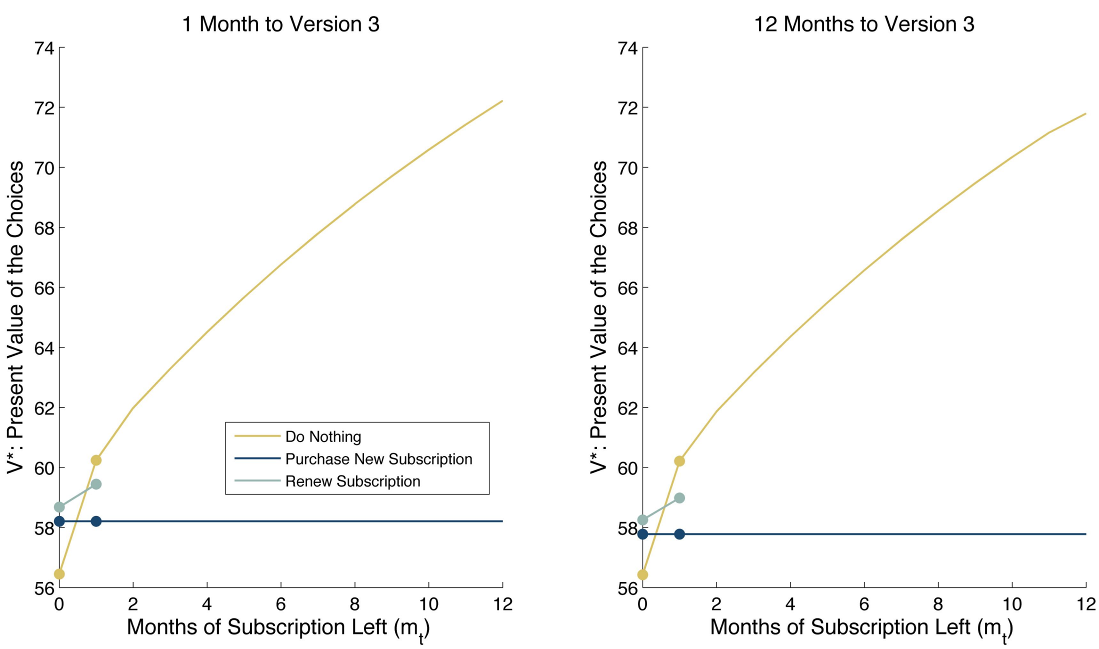


Parameter Values: Heterogeneous Model

## Figure 23: Type A Choice Specific Value Functions

at $\bar{e} = 2$, $(p_n, p_r) = (6, 8)$, Upgrade Available

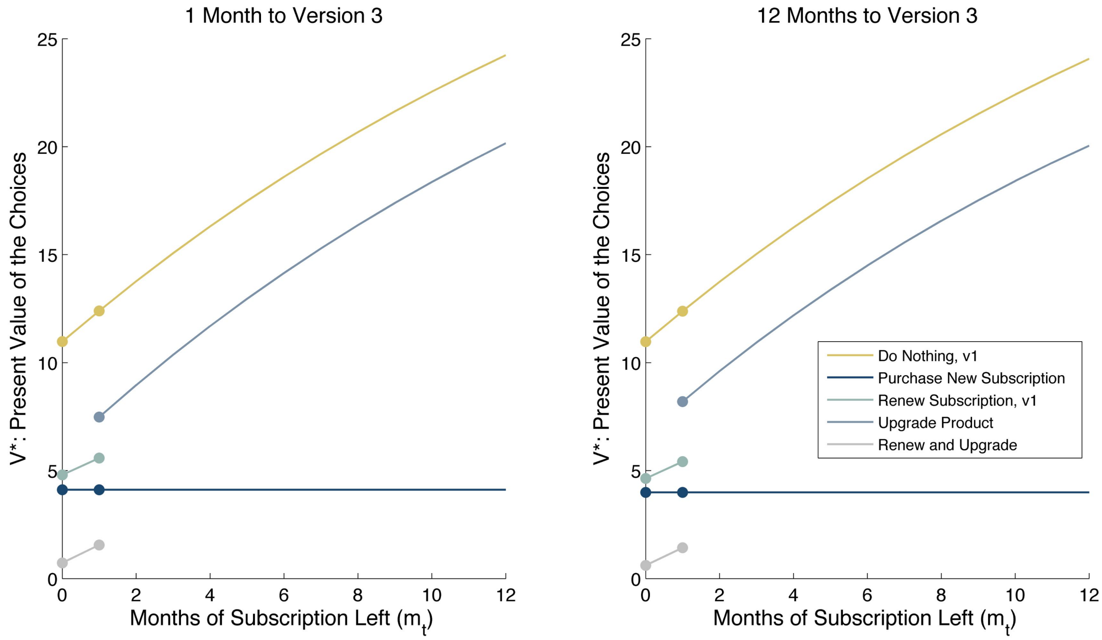


Parameter Values: Heterogeneous Model

## Figure 24: Type B Choice Specific Value Functions

at $\bar{e} = 2$, $(p_n, p_r) = (6, 8)$, Upgrade Available

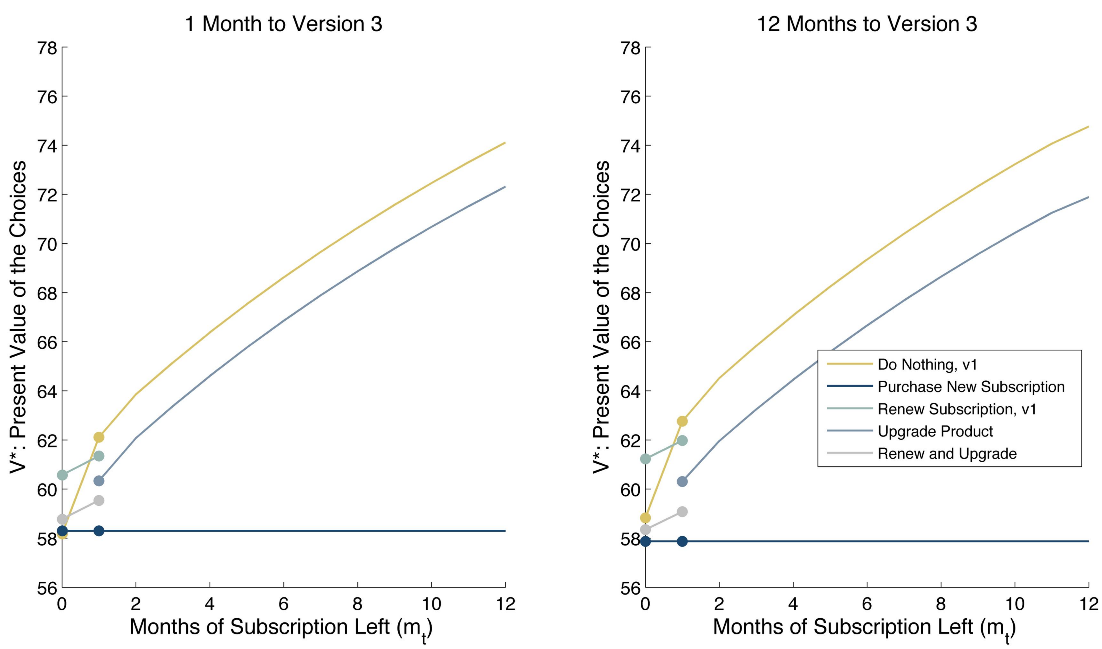


Parameter Values: Heterogeneous Model

# H Illustrative Model of the Firm's Optimal Pricing with Heterogeneous Version Valuation

In order to better understand how the firm's price discrimination decisions can be driven by the heterogeneity in consumers' valuations of successive versions, I examine a simple the firm's profit maximization problem in a world with two product versions and two types of consumers, who value each successive version differently. Each period, the firm allows the consumer to purchase a subscription to either of the product versions.

In this model, letting the consumer $i$'s' valuation for the product version $e$ be $\theta_e^i$, I can write the consumers' per-period utility as $(\theta_e^i - p_e)x$, where $x(\theta_e^i) \in \{0,1\}$ indicates whether the consumer $i$ chooses to purchase the product version $e$ and $p_e$ is the price of the product version. Allowing two types of consumers (type $H$ and type $L$) and two versions of the product (version 1 and version 2), I then have four possible valuations for the product, depending on the consumer type and the version.

$$\theta \in \{\theta_1^L, \theta_2^L, \theta_1^H, \theta_2^H\}, \quad \Pr\{\theta_e = \theta_e^H\} = \alpha, \ \forall e \tag{37}$$

In addition, assume for what follows that $\theta_e^L < \theta_e^H$, $\forall e$.

The firm then maximizes its profit with respect to prices (assuming zero marginal cost of production), subject to the incentive compatibility and individual rationality constraints of the consumers:

$$\max_{p_1, p_2} \pi = \alpha[x(\theta_1^H)p_1 + x(\theta_2^H)p_2] + (1-\alpha)[x(\theta_1^L)p_1 + x(\theta_2^L)p_2] \tag{38}$$

For the firm, there are two possible optimal actions:

1. Setting the price such that only the high type buys: the high type will buy the most appealing version
2. Setting the price such that both types buy: the types will sort according to their valuations of the two versions

**Case 1:** $\theta_2^H - \theta_1^H < \theta_2^L - \theta_1^L, \quad \theta_2^H < \theta_1^H, \quad \theta_2^L > \theta_1^L$ In this case, if the firm wants only the high types to purchase, it will set the price such that the high type buys version 1. Alternatively, if the firm wants both the low and the high type to purchase, the firm will set the price such that the low type sorts into purchasing version 2, while the high type purchases version 1.

$$\pi(p_1, p_2) = \begin{cases} \alpha p_1 & \text{if } p_1 = \theta_1^H, p_2 > \theta_2^H \\ \alpha p_1 + (1-\alpha)p_2 & \text{if } p_1 = \theta_2^L + (\theta_1^H - \theta_2^H), p_2 = \theta_2^L \end{cases} \tag{39}$$

The firm will choose to sell to high types only if the proportion of high types $\alpha$ is high enough such that the profits from selling to high types will exceed the profits of selling to both.

$$p^*(\alpha) = \begin{cases} p_1^* = \theta_1^H, p_2^* > \theta_2^H & \text{if } \alpha > \frac{\theta_2^L}{\theta_2^H} \\ p_1^* = \theta_2^L + (\theta_1^H - \theta_2^H), p_2^* = \theta_2^L & \text{if } \alpha < \frac{\theta_2^L}{\theta_2^H} \end{cases} \tag{40}$$

Given the parameter estimates in my data example, $\alpha = 0.26$, $\frac{\theta_2^L}{\theta_2^H} = 0.41$. Thus, $p_1^* = \theta_2^L + (\theta_1^H - \theta_2^H)$, $p_2^* = \theta_2^L$. Note, that in this case, it is optimal for the firm to sell version 1 at a higher price than version 2, and allow low type consumers to sort into version 1 and high type consumers to sort into version 2. We see this happening in my data application to some extent. Renewals of old versions are priced higher than first purchases of new versions, and the consumers renewing old versions are the high type consumers, while the low type consumers repurchase licenses of the newest version of the product.

Figure 25: Case 1 Prices (Given the Parameter Estimates)

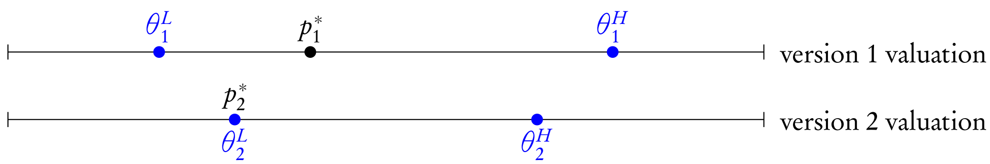


**Case 2:** $\theta_2^H - \theta_1^H > \theta_2^L - \theta_1^L$, $\theta_2^H > \theta_1^H$, $\theta_2^L < \theta_1^L$ Taking the opposite example, where the low types value the first version more than they value the second version, while the high types value version 2 more than version 1, we obtain the opposite result - one where the firm prices version 2 at a higher price than version 1 and the high types sort into version 2, while the low types sort into version 1.

I first consider the firm's profit, depending on whether it chooses to sell to high types only or to both high and low types, in this case.

$$\pi(p_1, p_2) = \begin{cases} \alpha p_2 & \text{if } p_1 > \theta_1^H, p_2 = \theta_2^H \\ \alpha p_2 + (1-\alpha) p_1 & \text{if } p_1 = \theta_1^L, p_2 = \theta_1^L + (\theta_2^H - \theta_1^H) \end{cases} \tag{41}$$

The firm will choose to sell to high types only if the proportion of high types $\alpha$ is high enough such that the profits from selling to high types will exceed the profits of selling to both.

$$p^*(\alpha) = \begin{cases} p_1^* > \theta_1^H, p_2^* = \theta_2^H & \text{if } \alpha > \frac{\theta_1^L}{\theta_1^H} \\ p_1^* = \theta_1^L, p_2^* = \theta_1^L + (\theta_2^H - \theta_1^H) & \text{if } \alpha < \frac{\theta_1^L}{\theta_1^H} \end{cases} \tag{42}$$

Taking the parameter estimates from my data and switching the estimated values for $\theta_1^H$ and $\theta_2^H$ (to allow increasing valuation for high types) and $\theta_1^L$ and $\theta_2^L$ (to allow decreasing valuation for low types), $\alpha = 0.26$, $\frac{\theta_1^L}{\theta_1^H} = 0.41$. Thus, $p_1^* = \theta_1^L, p_2^* = \theta_1^L + (\theta_2^H - \theta_1^H)$.

Figure 26: Case 2 Prices (Given the Parameter Estimates)

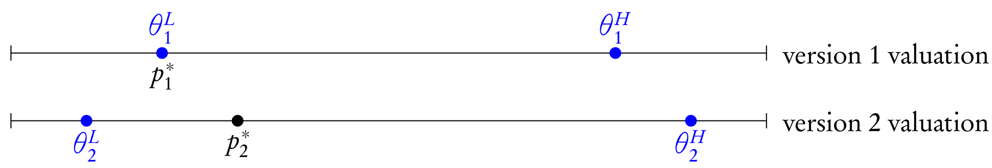


In this case, the new version is priced at a premium and the high types sort into buying the new version while the low types continue to purchase the old version at a lower price - this is the pricing often observed in various durable goods markets; e.g., cameras, TVs, etc.

**Case 3:** $\theta_2^H - \theta_1^H > \theta_2^L - \theta_1^L, \quad \theta_2^H > \theta_1^H, \quad \theta_2^L > \theta_1^L$ In this case, assuming the firm can convince the low types to sort into version 2, when indifferent, the firm will set the prices at the valuation of the low type for both versions, and everyone will continue to buy version 2.

I first consider the firm's profit, depending on whether it chooses to sell to high types only or to both high and low types, in this case.

$$\pi(p_1, p_2) = \begin{cases} \alpha p_2 & \text{if } p_1 > \theta_1^H, p_2 = \theta_2^H \\ p_2 & \text{if } p_1 = \theta_1^L, p_2 = \theta_2^L \end{cases} \tag{43}$$

The firm will choose to sell to high types only if the proportion of high types $\alpha$ is high enough such that the profits from selling to high types will exceed the profits of selling to both.

$$p^*(\alpha) = \begin{cases} p_1^* > \theta_1^H, p_2^* = \theta_2^H & \text{if } \alpha > \frac{\theta_2^L}{\theta_2^H} \\ p_1^* = \theta_1^L, p_2^* = \theta_2^L & \text{if } \alpha < \frac{\theta_2^L}{\theta_2^H} \end{cases} \tag{44}$$

Taking the parameter estimates from my data and switching the estimated values for $\theta_1^H$ and $\theta_2^H$ (to allow increasing valuation for high types), $\alpha = 0.26$, $\frac{\theta_2^L}{\theta_2^H} = 0.39$. Thus, $p_1^* = \theta_1^L, p_2^* = \theta_2^L$.

Figure 27: Case 3 Prices (Given the Parameter Estimates)

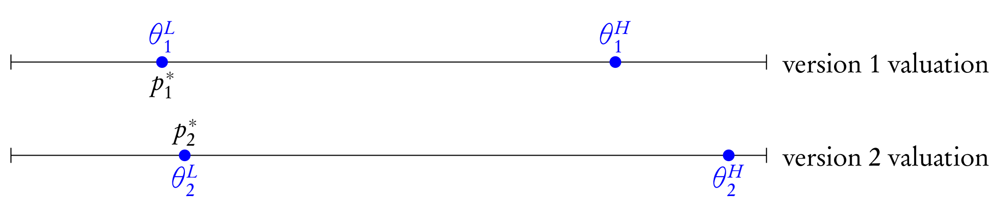


In this case, the new version has a slightly higher price, but there is no sorting across versions, and the two types both buy version 2.

**Case 4:** $\theta_2^H - \theta_1^H < \theta_2^L - \theta_1^L, \quad \theta_2^H > \theta_1^H, \quad \theta_2^L > \theta_1^L$ In this case, the firm prices version 1 high enough such that both types purchase version 2. I first consider the firm's profit, depending on whether it chooses to sell to high types only or to both high and low types, in this case.

$$\pi(p_1, p_2) = \begin{cases} \alpha p_2 & \text{if } p_1 > \theta_1^H, p_2 = \theta_2^H \\ p_2 & \text{if } p_1 > \theta_1^L + (\theta_2^H - \theta_1^H), p_2 = \theta_2^L \end{cases} \quad (45)$$

The firm will choose to sell to high types only if the proportion of high types $\alpha$ is high enough such that the profits from selling to high types will exceed the profits of selling to both.

$$p^*(\alpha) = \begin{cases} p_1^* > \theta_1^H, p_2^* = \theta_2^H & \text{if } \alpha > \frac{\theta_2^L}{\theta_2^H} \\ p_1^* > \theta_1^L + (\theta_2^H - \theta_1^H), p_2^* = \theta_2^L & \text{if } \alpha < \frac{\theta_2^L}{\theta_2^H} \end{cases} \quad (46)$$

Taking the parameter estimates from my data and switching the estimated values for $\theta_1^H$ and $\theta_2^H$ (to allow increasing valuation for high types) and allowing $\hat{\theta}_2^L - \hat{\theta}_1^L = 0.3$, $\alpha = 0.26$, $\frac{\theta_2^L}{\theta_2^H} = 0.44$. Thus, $p_1^* = \theta_1^H, p_2^* = \theta_2^L$.

Figure 28: Case 4 Prices (Given the Parameter Estimates)

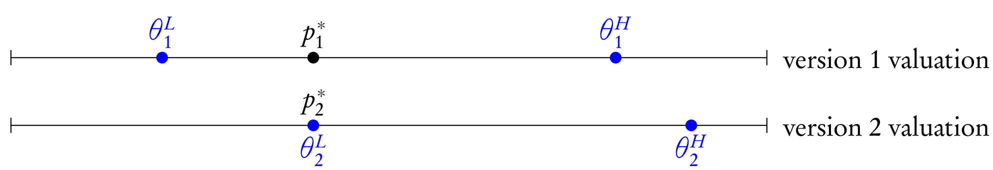


In this case, the old version has the same price than the new version (or slightly lower or higher), but there is no sorting across versions, and the two types both buy version 2.