

Welfare Analysis of Equilibria With and Without Early Termination Fees in the US Wireless Industry

by Joseph Cullen,¹ Nicolas Schutz² and Oleksandr Shcherbakov³

¹ Amazon

jacullen@gmail.com

² Department of Economics
University of Mannheim

nicolas.schutz@gmail.com

³ Currency Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

oshcherbakov@bankofcanada.ca



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

We would like to thank Victor Aguirregabiria, Philipp Schmidt-Dengler, Gautam Gowrisankaran, Kim P. Huynh, Yuya Takahashi, and participants of various seminars and presentations. We acknowledge the use of the Bank of Canada High Performance Cluster EDITH2. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada. All remaining errors are the responsibility of the authors.

Abstract

We study the social welfare implications of early termination fees in the US wireless industry. It is hypothesized that the elimination of long-term contracts at the end of 2015 was a transition from one market equilibrium to another. We use a theoretical model to illustrate that the endogenous choice of consumer switching costs by service providers does not necessarily raise firms' profits or hurt consumers. The forward-looking behavior of consumers facing switching costs results in significant downward pressure on prices. Service fees may be so low that consumers are better off and firms are worse off in an equilibrium with switching costs. Empirically, we find that without early termination fees, firms would increase prices by 2 to 5 percent, on average, resulting in an unambiguous increase in consumer surplus. Firms' profits derived from monthly service fees also increase. However, if we consider additional revenues from contract termination payments, the cost of processing these payments should be large enough for producer profits to be higher in the new equilibrium.

Topics: Econometric and statistical methods; Firm dynamics; Market structure and pricing

JEL codes: D22, L15, L96

1 Introduction

In a competitive environment, acquiring and retaining customers is the central objective of firms. Firms want to attract customers from competitors while protecting their own customer base. One defensive strategy a firm may employ is to create exit barriers for existing customers. Increasing the cost of switching to another product or service makes it more difficult for rival firms to poach customers and tends to lessen competitive pressures. Firms may introduce strategic incompatibilities, create artificial network effects, or explicitly write contracts that embody switching costs. Firm strategies that lock consumers in could harm consumers and reduce welfare. Klemperer (1995) argues that switching costs may result in substantial welfare losses and suggests that social planners work to reduce them. However, more recent theoretical and empirical work on switching costs emphasizes that consumers may benefit from the presence of switching costs (Dube et al., 2009). That is, even though creating switching costs could be a profitable unilateral strategy, it may backfire when used by all the firms in the market. The difficulty of acquiring customers that have been locked in by competitors may outweigh the benefits of preventing your own customers from switching when consumers are forward looking. In this paper we empirically examine welfare in an industry where firms endogenously choose whether or not to have switching costs.

The setting for our study is the US wireless phone market. In 2014, revenues in the mobile telephony market exceeded 200 billion dollars per year representing 355 million unique subscribers (CTIA, 2015).¹ Wireless connectivity has become increasingly important in day-to-day living. In 2014, more than 45 percent of US households relied on a wireless phone as their only phone line (Blumberg and Luke, 2015). It is also increasingly common for wireless service to be an individual's primary connection to the internet (Pew Research Center, 2015). Due to its importance, concerns about the competitiveness of the market are paramount. There are four large, national wireless providers in the US. In 2011, a major merger between the second and fourth largest wireless providers was blocked by the US Department of Justice due to the potential anti-competitive effects. However, as consumer advocates have pointed out, having many firms in the market may not make much of a difference for consumers if the cost of switching providers curtails competition between providers.

Switching costs are explicitly embodied in most wireless service contracts through early termination fees (ETFs). The vast majority of contracts are multi-year (usually two-year) agreements for wireless service. ETFs penalize consumers for leaving the wireless provider before the end of their contract. Providers argue that ETFs are necessary due to phone subsidies. Signing a long-term contract with a firm allows the consumer to buy a handset at a subsidized price. The amount of the subsidy can be significant, in the hundreds of dollars for advanced devices like smart phones. Over the lifetime of the contract, the firm generates its profit by receiving monthly service fees paid by the customer. Firms argue that they only use ETFs to prevent consumers from buying a phone at a subsidized rate and then immediately canceling their contract, leaving the firm with no opportunity to recoup its subsidy through monthly service fees. However, the structure of ETFs suggest that they are also being used strategically by the firms to prevent switching. Since each month the firm recovers a part of the cost of subsidizing the handset, the early termination fee decreases over time. However, in practice ETFs decrease so slowly that there is a

¹In 2014, revenues from services fees were 187 billion dollars. Revenues from wireless services and equipment were 224 billion dollars.

substantial termination fee even in the final month of the contract. This suggests that firms may be choosing ETFs to prevent customers from switching to rivals.

It is safe to say that ETFs are widely unpopular with consumers. As such, they have also come under scrutiny by federal regulators and legislators as well as become the subject of several high-profile, class-action lawsuits². However, it isn't clear that consumers are worse off with ETFs than without ETFs. That is, regulation banning ETFs could actually harm consumers rather than benefit them if changes in the competitive structure lead to firms raising prices in the absence of ETFs.

Furthermore, it isn't clear that producers are better off with ETFs. This may seem counterintuitive. Since firms are choosing to have ETFs, shouldn't it be the case that they are better off with ETFs? However, this might not be the case if there are multiple equilibria. Since there is no theory to appeal to on endogenous switching costs, we investigate this question by developing a simple, theoretical model where firms endogenously choose whether to have ETFs and then compete in per-period prices. Specifically we seek to understand whether multiple equilibria could exist in our setting. We find that, in equilibrium, either all firms will have ETFs or none of the firms will have them. The theory suggests that consumers could actually be better off in equilibrium *with* switching costs, while firms might prefer the equilibrium *without* the ETFs. This implies that an ETF equilibrium could be the result of a coordination failure among firms rather than the firms' preferred outcome.

Having established theoretically that the effect of endogenous ETFs on consumers and producers is ambiguous, we continue our analysis by empirically investigating the effects of ETFs on welfare and competition. Using a detailed consumer survey from 2005 to 2012, we estimate a model of dynamic, forward-looking consumers in the market for wireless services. The data contain individual-level information on the purchase decisions, consumer demographic characteristics, and handset characteristics. The demand model is a literal application of the Berry et al. (1995) (BLP) model to a dynamic setting. Just like BLP assumes that consumers have perfect information about product characteristics in a static setting, our model assumes that consumers know the characteristics of all the products in the market both now and in the future. This perfect foresight approach to consumer dynamics differs from prior work in the dynamic demand literature, which generally assumes some parametric form for uncertainty in consumer beliefs over the future. Although perfect foresight rules out uncertainty in consumer beliefs, it endows distinct advantages to the model.

First, it allows consumers to account for the evolution of each product separately, which we believe is important to accurately model the decision-making process in the wireless market. This contrasts with Markovian models of consumer uncertainty, which typically model beliefs about the future using fairly strong assumptions on sufficient statistics (e.g., values of the current holding and logit inclusive value) and parametric restrictions on the evolution of these aggregates.³ For example, in our model, a consumer can believe that two products, which currently have the same current period utility, will evolve very differently in the future. Second, by using the actual product evolution for each product, we can avoid approximation errors in the estimation of beliefs. With perfect foresight, it is clear that the actual evolution of product characteristics is driving the results. Finally, our demand-side model is entirely consistent with our supply-side

²https://apps.fcc.gov/edocs_public/attachmatch/DOC-295965A1.pdf

³Examples of applications can be found in Schiraldi (2011), Gowrisankaran and Rysman (2012) and Nosal (2012). Aguirregabiria and Nevo (2013) provide a detailed discussion of the issue.

approach. Modeling producer decisions with forward-looking consumers facing uncertainty on the demand side is computationally challenging. In our model, we assume that the evolution of characteristics is common knowledge for both consumers and producers, with the latter employing open-loop strategies.

Our structural model accounts for two important sources of dynamics: handset durability and switching costs (ETFs). In addition, we allow for sticky contract prices when the level of monthly service fees is fixed in the long-term contract. Importantly, we don't estimate the "hassle" costs associated with switching providers, but instead focus on the explicit switching costs set by firms in wireless service contracts by including data on actual ETFs in the model. Because there is no extra, hassle cost of switching to estimate, identification in our model is more transparent. Focusing only on ETFs can be viewed as estimating a lower bound of switching costs. It is worth noting that prior work, which does estimate total consumer switching costs in the industry, finds that the estimates are very close to the average ETF in service contracts (see Cullen and Shcherbakov, 2010).

Individual-level data allow us to accommodate some heterogeneity in the estimation procedure by allowing estimates to vary flexibly across demographic groups. Note the dynamic model is key for producing sensible counterfactual results. If consumers were completely myopic, firms could use ETFs to lock in consumers permanently. Dynamic, forward-looking consumers curb the ability of firms to extract rents through future switching costs in counterfactual simulations.

The final step of the paper is to estimate supply-side structural parameters and then simulate several counterfactual scenarios to empirically investigate whether firms or consumers are better or worse off without ETFs. The results are particularly interesting, as the industry recently transitioned to a business model without ETFs and long-term contracts. We provide a set of counterfactual simulations within both partial and full equilibrium analysis. Our partial equilibrium analysis suggests that the elimination of the early termination fees would increase consumer welfare by about 76 percent if handset prices and service fees remained unchanged. This is only a partial equilibrium result since wireless providers would undoubtedly change service fees and handset prices without an ETF. In particular, firms would probably not sell subsidized handsets without an ETF. However, even if consumers face the full, unsubsidized prices for handsets, the increase in consumer value functions is still estimated to be at about 48 percent when monthly service fees remain unchanged.

Within our partial equilibrium framework, we also examine the relative importance of dynamics due to the handset durability vs. consumer switching costs. In particular, if we remove handset durability by allowing consumers to rent, rather than buy, handsets, consumer welfare increases by 19 percent.⁴ When we eliminate both sources of dynamics by allowing consumers to rent their handsets on a per-period basis and eliminate the early termination fees, consumer welfare increases by about 116 percent on average.

These partial equilibrium results don't account for changes in services fees that firms may implement in the absence of ETFs. In a new no-ETF equilibrium where consumers face unsubsidized phone prices, a 32 percent proportional increase in service fees would, on average, eliminate the consumer welfare increase due to the elimination of early termination fees. Therefore, in equilibrium, where prices increase significantly, the elimination of the early termination fees need not be beneficial to consumers because a sufficiently large increase in prices may offset the positive effects of the ETF elimination.

⁴The per-period rental price for handsets is set equal to the average handset depreciation rate of 24 percent per six months.

To simulate a full equilibrium scenario where the wireless providers optimally choose their service fees in a situation without ETFs, we first recover their cost structure. We assume that the service providers set their prices to maximize joint profit from all of their products in all time periods.⁵ Costs of providing each product are then recovered in a way similar to Berry et al. (1995) using a set of first-order conditions. Full equilibrium is then simulated under the assumption of no ETFs but with consumer decisions still affected by the durability of the handsets.

In the new equilibrium without ETFs, service fees would be higher by 2.10 to 5.17 percent on average with larger increases for bigger carriers. Consumer surplus is higher by 68 percent such that the lifetime value of the wireless service for consumers increases from 2,347 to 3,932 dollars. Producer profits from service fees also increase by 46 to 89 percent, with smaller carriers gaining more. However, elimination of ETFs also eliminates revenues received from these payments. If the costs of processing ETFs were zero, carrier profits would be smaller without ETFs. However, if the per-ETF processing costs were high enough, e.g., constituting at least 8.39 dollars for Verizon, 6.51 dollars for AT&T, 2.81 dollars for Sprint, and less than 2 dollars for smaller carriers, then producer profits would be higher in an equilibrium without ETFs and long-term contracts. We rationalize the difference between the theory model results predicting larger changes in equilibrium prices and our finding of a relatively small price increase by correlation of quality for competing products over time. The theoretical setup assumes perfect negative correlation in qualities of two products in the market. Therefore, streams of flow utilities generated by the alternatives over an infinite time horizon look very similar from the consumer point of view. This homogeneity over time is eliminated if consumers can freely choose any product every period. Our empirical results suggest a very strong positive correlation in service fees of competing providers. Therefore, a smaller-than-theoretically-predicted increase in prices makes sense.

The rest of the paper is organized as follows. In Section 2 we discuss the role of switching costs using a simple theoretical model that illustrates potential multiplicity of equilibria in a wireless market and motivates our further empirical analysis. Section 3 describes our data. The dynamic model of consumer demand and estimation algorithm are discussed in Section 4. We provide the identification argument and discuss instrumental variables in Section 5. Estimation results are reported in Section 6. We provide extensive robustness checks for assumptions on the initial conditions and consumer beliefs in the terminal period. We also conduct robustness exercises with respect to the set of instrumental variables used in estimation. Section 7 outlines the model of wireless operators and presents results of our counterfactual simulations. Here, we scrutinize additional assumptions made for computational tractability of the supply side in our model. Section 8 concludes.

2 Theoretical predictions

2.1 The model

We begin by introducing an illustrative, theoretical model of competition with switching costs. Since we only see an equilibrium *with* ETFs in the data, one might wonder whether it is a dominant strategy to impose switching costs on customers in this context. If this were

⁵Under the assumption of open-loop strategies with consumers having perfect foresight over the evolution of individual products, the firms take into account the fact that a change in price of any product affects the demand for all its products in all time periods.

the case, then ETFs would disappear only if they were banned by a regulatory authority. On the contrary, our theoretical investigation finds that there can exist equilibria in which firms choose not to impose switching costs on consumers without government intervention. In addition, for economically sensible parameters, the theoretical model predicts that both the ETF and the no-ETF equilibria exist. This suggests a possibility of coordination problems.

We provide a brief overview of the theory framework and predictions here. The details of the model and formal proofs of the results can be found in Appendix A. The framework is a Hotelling model with two firms and a continuum of consumer types. Time is discrete and runs from $t = 0$ to $t = \infty$. Before the dynamic game starts, firms 1 and 2 simultaneously and non-cooperatively decide whether to use ETFs. This commitment is made once and for all. Next, the dynamic pricing game starts. In each period, firms choose the price for their products and consumers realize their type before choosing which firm to buy from. If ETFs are present, consumers take into consideration that they will be locked into purchasing from the same firm in the future, even though their type will be evolving over time. Consumers are risk neutral and discount the future at rate $\beta_c \in [0, 1)$. Likewise, firms are risk neutral and discount the future at rate $\beta_f \in [0, 1)$. We look for stationary subgame-perfect equilibria in pure strategies.

Let (e, e) and (n, n) denote two symmetric equilibria when both firms choose to introduce ETFs or both of them choose not to introduce ETFs, respectively. Similarly, (e, n) and (n, e) denote asymmetric equilibria when only one of the firms choose to have ETFs. Formal representation of the game can be found in Appendix A. Its solution is given by two cut-off functions $\beta_c^n(\beta_f)$ and $\beta_c^e(\beta_f)$. Our main theoretical result is presented in Proposition 1 below.

Proposition 1. *Let $(\beta_c, \beta_f) \in [0, 1)^2$. Then:*


- (n, n) is an equilibrium if and only if $\beta_c \geq \beta_c^n(\beta_f)$.
- (e, e) is an equilibrium if and only if $\beta_c \leq \beta_c^e(\beta_f)$.
- (n, e) and (e, n) are equilibria if and only if $\beta_f = \beta_c = 0$.

For every $\beta_f > 0$, the set of β_c s such that equilibria (n, n) and (e, e) coexist is an interval with a non-empty interior.

Proof. This follows immediately from Lemma 4; see Appendix A. □

The proposition states that the equilibrium of the model depends on the relative patience of producers and consumers. Figure 1 provides a graphical representation of Proposition 1 in the (β_f, β_c) plane. The lightly shaded area in the graph shows the range of parameters where only an ETF equilibrium can exist. It is a dominant strategy for firms to unilaterally impose an ETF even if their competitor does not choose an ETF. This only occurs if firms are much more patient than consumers. The unshaded area in the top left of the graph indicates the range of discount factors where in equilibrium there would be no ETF. This usually occurs when consumers are more patient than producers.

Figure 1: Equilibrium characterization



The most interesting and economically relevant outcome occurs in the darkly shaded area. Here there are exactly two equilibria. Either both firms impose an ETF or neither firm does. The range of discount factors where this occurs is also economically relevant. If consumers are less patient than the firms, then multiple equilibria are likely to exist and firms face a coordination problem.

Importantly, the firms and consumers are not indifferent between the two equilibria. On the one hand, the ETF equilibrium yields lower payoffs to firms and worse overall allocative efficiency than the no-ETF equilibrium (see Appendix A.2.2). On the other hand, consumers are better off under the ETF equilibrium due to much lower prices. The intuition behind these findings is related to consumer forward-looking behavior and product differentiation. In our theoretical setup, quality of the competing products varies a lot and is strongly negatively correlated with each other. Without ETF such a market structure would result in a sequence of static differentiated Bertrand markets where differences in product qualities map into differences in prices (here we ignore for simplicity other sources of dynamics such as handset durability and sticky contract prices). With infinite ETFs consumers never switch. They have to choose between two streams of flow utilities correlated negatively over time. Note that despite heterogeneous qualities in every given period these differences are smoothed over time such that both options look very similar from today's perspective. As products become more homogeneous from the consumer point of view, firms' ability to set high prices declines. Therefore, the difference in equilibrium prices between ETF and no-ETF scenarios depends on the difference in product heterogeneity when they are sold in bundles or à la carte over time. In the wireless industry, we expect product quality to be strongly positively correlated over time and, hence, we do not expect a significant increase in equilibrium service fees after long-term contracts and ETFs are eliminated.

In the empirical application, we relax several strong assumptions made in the theoretical model. In particular, we allow for several multi-product service providers who set finite early termination fees that depend on the contract time and the type of handset. In our model, consumers sometimes switch and pay ETFs. They can also choose an outside

option represented by the best non-contract telecommunication service available. We also allow firms' profits to depend on ETF payments. Using data on US cellular customers, we estimate consumer preferences and firm costs. In this way, we assess the effect of early termination fees on the equilibrium prices of service providers and the resulting consumer and producer surplus.

3 Data

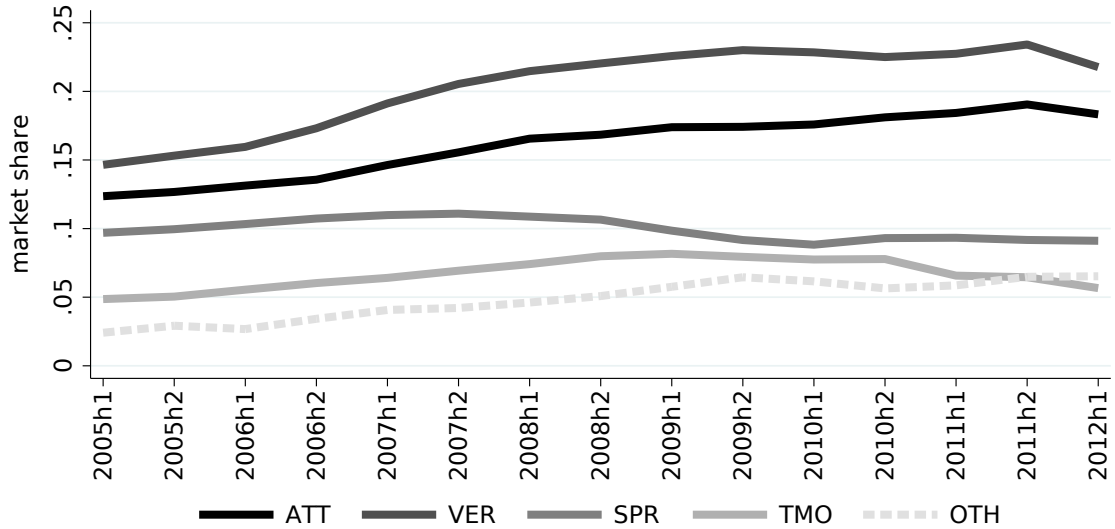
We use data from a quarterly cross-sectional consumer survey collected by *Comscore Inc.* from 2005 to 2012. *Comscore* administers the detailed survey to a random sample of approximately 36,000 cell phone users each quarter to quantify market growth and cell phone usage patterns. The survey includes questions on handset used, price paid for the handset, current carrier, monthly fee for the calling plan, demographic characteristics of the individual, and other factors. In addition, *Comscore* maintains a database of detailed handset characteristics that can be matched to the cell phone model owned by an individual. The sample of consumers is weighted and balanced to match national subscriber numbers and demographic characteristics.

The major wireless service operators include Verizon, AT&T, T-Mobile, and Sprint, all of which offer virtually nationwide service. We aggregated all other regional or local wireless carriers into a separate category labeled "Other."⁶ The top four carriers account for the vast majority of cell phone users; approximately 90% of the users in the sample subscribed with one of these four major carriers.

Our product definition is a handset-carrier combination. For example, we define the iPhone 4S on the AT&T network as a single product and calculate its market share for each year in the sample as the total number of projected subscribers divided by the US population. The total market share of the carrier is simply the sum of the market shares of each of its products. Figure 2 shows how the market shares of carriers evolved over the course of the sample. On average market shares of carriers have increased over time as cell phones have become increasingly common. Note that carrier market shares mask the rich variation in handset-carrier market shares. The latter facilitates identification in our model.


⁶There have been a few significant mergers in the industry. The largest of these was the merger of Cingular and AT&T, which occurred before the beginning of our sample. A smaller, but still sizable, merger occurred in 2009 when Verizon acquired Alltel wireless, which at the time was the fifth largest wireless company. The data provided by Comscore retroactively aggregated the market shares of Alltel and Verizon together for the whole sample. Thus Verizon's market shares represent the combined market shares of Alltel and Verizon customers before 2009.

Figure 2: Market share by carrier-year, 2005-2012



The survey also includes information on monthly service fees and, for those who purchased a phone in the current month, the price they paid for their handset. For the price of the handset market, we use the average reported handset price by individuals in that year. For the carrier monthly fee, we use the average monthly fee for all subscribers to that carrier. Figure 3 illustrates the evolution of average handset prices (left panel) and service fees (right panel) as reported by the survey participants.

Figure 3: Average handset prices (left) and service fees (right) by carrier-year, 2005-2012




Note: Reported handset prices are weighted by the number of respondents.

Since handsets are durable goods, the number of possible handsets on each carrier increases over time as new handsets are introduced by each carrier annually; when estimating the structural model we assume that any handset available in earlier years could be used in later periods due to its durability. It is worth noting that the survey may not contain information on market shares for all possible handset-carrier combinations. Therefore, while our model will predict the entire distribution of shares, to form moment conditions we match the model predictions only to the observations available from the

survey. Figure 4 summarizes the average number of distinct handset models used with each of the main wireless service providers.

Figure 4: Number of handsets by carrier-year, 2005-2012



In Table 1 we list some of the handset characteristics available in our data. Although we estimate our dynamic model using handset dummy variables, detailed handset characteristics are used to construct instrumental variables for the estimation based on product “similarity,” as discussed in Section 5.

To estimate our model, we aggregated quarterly data to the bi-annual level. The aggregation is used to obtain more precise measures of market shares at the handset-carrier-time level. Our estimation algorithm relies on a dynamic version of the inversion method originally proposed by Berry et al. (1995). Therefore, an accurate measure of the population purchase probabilities is important for the consistency of our estimates.


Table 1: Selected handset characteristics

variable name	variable name
Smartphone (y/n)	GPS (y/n)
Built-in storage (y/n)	Email (y/n)
Java version (MIDP 2.0, Dalvik, etc.)	Full-keyboard (y/n)
Bluetooth (y/n)	GPRS (y/n)
Infrared (y/n)	IM (y/n)
Display width	MMS (y/n)
Display height	MPEG-4 (y/n)
Display color (65,536; B&W, etc.)	Formfactor (Candybar, Slider, etc.)
Audio type (realtones, monophonic, etc.)	Release date (year/q)
GSM (y/n)	OS type (Microsoft, Symbian, etc.)
CDMA (y/n)	Camera resolution (mgp)

Notes: GSM stands for global system for mobile communications; CDMA for code-division multiple access; GPS for global positioning system; GPRS for general packet radio services; IM for instant message; MMS for multimedia messaging service; and MPEG for moving picture experts group standards.

We collect data on carrier-specific early termination fees from carrier websites and past announcements. The ETF schedules are generally fixed for years at a time. We only observe one revision of ETFs for each carrier over the sample. The shift for each carrier seems to be motivated primarily by the cost of smart phones. Over the 2009-2011 time period each carrier introduced a higher ETF for “advanced devices” (smart phones), which generally enjoyed a higher upfront carrier subsidy. Even though advanced ETFs are much larger, the percentage decline over the lifetime of the contract is very similar to the ETFs for basic devices. Figure 5 shows the ETFs for advanced devices for each month of a two-year service contract.

Figure 5: Early termination fees for advanced devices by carrier



There are several caveats related to the survey data we use for estimation. We observe only expenditures and not the exact characteristics of the service plans. Therefore, in the estimation, we assume that all wireless subscribers choose the same service plan. This is not an innocuous assumption because the same model of handsets can be offered with very different plans (e.g., large data plan, family plan, etc.). Our sample begins in 2005 and ends in 2012, which creates initial conditions and terminal period problems for our empirical dynamic model. We address both issues in Sections 4 and 5 below.

4 Dynamic demand

4.1 Consumer behavior

We define a product to be a particular handset-carrier combination. Handsets are assumed to be useless without a subscription to a wireless service provider, and there is no value in subscribing to the service without a hardware device. Time is discrete and corresponds to a six-month interval, consistent with the data we use for estimation. The set of products in each time period, t , is denoted by $\mathcal{J}_t \subseteq (\mathcal{H}_t \times \mathcal{C}) \cup \{o\}$, where \mathcal{H}_t is the set of cellular phones available in period t , \mathcal{C} is the set of wireless service providers, and o denotes the outside option of not using contract-based wireless communication services. We denote J_t the cardinality of set \mathcal{J}_t . Handsets are durable, do not depreciate, and can be used for many periods. We assume that consumers cannot choose a different service provider without buying a new handset.⁷

Each product is described by a vector of observable attributes $x_{jt} \in \mathbb{R}^K$, which includes handset features as well as characteristics of the service provider. For example, a handset is characterized by its price, display type and size, availability and resolution of camera, form factor, networking capabilities, etc. Service providers differ in terms of coverage quality, customer service, and monthly subscription fees. We also allow for product characteristics that are observed by consumers and firms, but unobserved by us. These characteristics are summarized by a scalar $\xi_{jt} \in \mathbb{R}$.

We assume that the market for wireless services is populated by a finite number of consumer types, $i = 1, \dots, N$. Each consumer type corresponds to a demographic group and is characterized by a finite-dimensional parameter vector ω_i , which summarizes the tastes for various product attributes. We assume that all contracts have a duration of two years, which we denote as $\mathcal{T} = 4$, the number of bi-annual periods. At the beginning of each time period, a consumer decides whether to sign a new four-period contract or to continue with the current service provider. Premature termination of a contract is costly. Let $d_{it} \in \mathcal{J}_t \cup \{\emptyset\}$ denote the consumer's purchase decision. Note that the set of feasible actions contains the possibility of not making an active purchase decision (i.e., not signing a new contract, but rather staying on the same contract as in the previous period), in which case $d_{it} = \emptyset$.

At the beginning of each period consumers are endowed with current holdings $e \in \bigcup_{t'=1}^{t-1} \mathcal{J}_{t'}$ defined as one of the products ever available in the market prior to the current period or the outside option. A pair (j, τ) uniquely identifies a consumer holding e by describing the identity of the product j purchased at one of the previous time periods τ . For example, a consumer can be endowed with an *iPhone 4* purchased under the contract with *AT&T* on March 24, 2011. In this case we denote this holding as $e =$

⁷This is true for the early years of our data but may be restrictive for the most recent years.

(iPhone 4 with AT&T, 03/24/2011). We assume that when a new product is purchased, the old endowment is disposed of at no cost. It is worth noting that in our model, the per period utility flows of a holding evolve over time (for example, due to changes in service quality).

Since we define $\tau \leq t$ to be a variable recording the time of the most recent purchase, $t - \tau$ determines the age of the consumer endowment. Therefore, a vector (e, t) completely describes the age and value of any endowment at the beginning of each period. Let $\alpha_{ip} \in \omega_i$ denote price sensitivity of a consumer type i . Then, an early termination fee of F_{et} and a handset price of P_{jt} have utility costs of $\eta_i(e, t) \equiv \mathbf{1}(t - \tau < \mathcal{T}) \alpha_{ip} F_{et}$ and $\gamma_{ijt} \equiv \alpha_{ip} P_{jt}$, respectively. Each contract is characterized by monthly subscription fees, p_e , which are specified at the beginning of the contract and stay constant over time unless the contract is terminated by the consumer. We denote $\phi_{ie} \equiv \alpha_{ip} p_e$ the utility cost of such per-period payments. Let $X_{jt} = (x_{jt}, \xi_{jt}, F_{jt}, P_{jt})$ and $X_t = (X_{jt})_{j \in \mathcal{J}_t}$. We assume that the consumer per-period utility function is given by

$$U_i(d, e, t, X_t, \bar{\varepsilon}_t) = \begin{cases} \delta^f(x_{et}, \xi_{et}, \omega_i) + \varepsilon_{iet} - \phi_{ie}, & \text{if } d = \emptyset, \\ \delta^f(x_{dt}, \xi_{dt}, \omega_i) + \varepsilon_{idt} - \phi_{i(d,t)} - \eta_i(e, t) - \gamma_{idt}, & \text{otherwise,} \end{cases} \quad (1)$$

where ε_{iet} and ε_{ijt} ($j \in \mathcal{J}_t$) are idiosyncratic match values given by i.i.d. random draws from a standard Gumbel distribution and $\bar{\varepsilon}_{it} = (\varepsilon_{iet}, (\varepsilon_{ijt})_{j \in \mathcal{J}_t})$. As discussed before, the consumer can either avoid making an active purchase decision and keep endowment e , or purchase any of the currently available products $d \in \mathcal{J}_t$. In the latter case, the consumer's current holding $e = (j, \tau)$ is replaced by $e' = (d, t)$.⁸

We assume that consumers maximize the expected present discounted value of future utility flows and that they have perfect foresight over future product attributes, except for the i.i.d. draws of $\bar{\varepsilon}_{it}$. Then, all state variables in the consumer dynamic programming problem can be summarized by $(e_i, t, \bar{\varepsilon}_{it})$. Let $\delta_{ijt}^f \equiv \delta^f(x_{jt}, \xi_{jt}, \omega_i)$. The consumer's dynamic problem can then be formulated recursively with the following Bellman equation (we omit subscript i from now on for notational convenience),

$$V(e, t, \bar{\varepsilon}) = \max \left\{ \begin{array}{l} \delta_{et}^f + \varepsilon_{et} - \phi_e + \beta \mathbb{E}V(e, t'), \\ \max_j \left[\delta_{jt}^f + \varepsilon_{jt} - \phi_{jt} - \eta(e, t) - \gamma_{jt} + \beta \mathbb{E}V((j, t), t') \right] \end{array} \right\}, \quad (2)$$

where

$$t' = \begin{cases} t + 1, & \text{if } t < T, \\ T & \text{otherwise,} \end{cases}$$

and

$$\mathbb{E}V((j, \tau), t) \equiv \int \cdots \int V((j, \tau), t, \varepsilon) dF(\bar{\varepsilon}).$$

We assume that after the terminal period in our data, T , consumers believe that the flow utility of each product evolves over time according to the following deterministic process:

$$\delta_{jt} = \hat{\gamma}_{0,g} + \hat{\gamma}_{1,g} \delta_{jt-1}, \quad (3)$$

where g denotes carrier identity, $\hat{\gamma}_{0,g}$ and $\hat{\gamma}_{1,g}$ are ordinary least squares (OLS) parameter estimates from the autoregressive model of order one (AR(1)),

$$\delta_{jt} = \gamma_{0,g} + \gamma_{1,g} \delta_{jt-1} + \nu_{jt}, \quad j \in \mathcal{J}_t, \quad t = 0, \dots, T,$$

⁸Note that we allow for repeated purchases of the same products.

and δ_{jt} are the current estimates of the mean flow utility. In Section 6, we provide robustness checks for this assumption. In particular, we consider an alternative setup, when consumers believe that flow utilities stay constant at their terminal period values, i.e., $\delta_t = \delta_T \forall t \geq T$. We also experiment with a partially stochastic evolution of the flow utilities.⁹

By using the standard properties of the Gumbel distribution we can rewrite the Bellman equation (2) in terms of $\mathbb{E}V(e, t)$,

$$\mathbb{E}V(e, t) = \ln \left(\begin{array}{l} \exp \left(\delta_{et}^f - \phi_{et} + \beta \mathbb{E}V(e, t') \right) \\ + \exp \left(\delta_{jt}^f - \phi_{jt} - \eta(e, t) - \gamma_{jt} + \beta \mathbb{E}V((j, t), t') \right) \end{array} \right). \quad (4)$$

To solve the dynamic programming problem for consumer type i , we proceed as follows. First, we augment our data with 60 periods after the terminal period T . We populate these periods with mean flow utilities defined in equation (3). For handset prices and service fees we also use auto-regressive specifications to estimate parameters and forecast future values using equations analogous to 3. Schedules of ETFs are assumed to remain the same for all periods after the terminal one.

We solve the consumer dynamic programming problem by using a backwards induction algorithm starting with the final period rewards identical to zero. We experiment with shorter and longer time spans left after the terminal period in our data and find that there is no difference in parameter estimates for all models extended by 60 or more time periods forward. There are only minor differences in parameters estimated using models extended by 30 and 60 periods forward.

It is clear that we never observe the δ_{jt}^f s in the data. However, we observe consumer purchase decisions, which can be mapped into the unobserved per-period utility values as we describe in the next section.

4.2 Computing market shares

Our estimation algorithm uses data on the product-level market shares, and we begin by defining them. Let s_{jt} (resp. s_{ijt}) denote the observed aggregate market share (resp. the observed market share in consumer group i) for a particular handset-carrier combination j at time t . Note that current-period purchase decisions depend on the current holdings, which are determined by the consumers' decisions in the previous periods. Due to the unbounded support of the vector of shocks, $\bar{\varepsilon}_{it}$, each consumer type purchases every product with positive probability. Therefore, to determine purchase probabilities for any product at time t , we need to circle over all products ever available in the market up to the current time period.

More formally, let $\Pr_i(d = j|e, t)$ denote the conditional probability that consumer type i buys product j at period t , given that the consumer's current holding is e . This

⁹For computational reasons, we don't experiment with changing the number of products after the terminal period and instead assume a constant number of products in all post- T periods.

probability is given by (subscript i omitted):

$$\begin{aligned} \Pr(d = j|e, t) &= \\ &\int \mathbb{1} \left(\begin{aligned} \delta_{jt}^f - C_{ejt} + \beta \mathbb{E}V((j, t), t') + \varepsilon_{jt} \geq \delta_{et}^f - \phi_e + \beta \mathbb{E}V(e, t') + \varepsilon_{et}, \\ \delta_{jt}^f - C_{ejt} + \beta \mathbb{E}V((j, t), t') + \varepsilon_{jt} \geq \delta_{kt}^f - C_{ekt} + \beta \mathbb{E}V((k, t), t') + \varepsilon_{kt}, \forall k \end{aligned} \right) dF(\bar{\varepsilon}_t) \\ &= \frac{\exp \left(\delta_{jt}^f - C_{ejt} + \beta \mathbb{E}V((j, t), t') \right)}{\exp \left(\delta_{et}^f - \phi_e + \beta \mathbb{E}V(e, t') \right) + \sum_{k \in \mathcal{J}_t} \exp \left(\delta_{kt}^f - C_{ekt} + \beta \mathbb{E}V((k, t), t') \right)}, \end{aligned} \quad (5)$$

where $C_{ejt} = \eta(e, t) + \gamma_{jt} + \phi_{jt}$.

To compute current period product shares for each consumer type i , we define a matrix of current consumer holdings, $H^i(t)$. The matrix has dimension $J_T \times T$, where J_T is the number of products in the final period of our data. Note that in our data \mathcal{J}_T contains all products ever available in the market. A generic element of matrix $H^i(t)$ is $h_{(j, \tau)}^i(t)$.

We assume that at the beginning of the market, $h_{j,0}^i(0) = 0$ for all i, j ; i.e., all consumers hold the outside option.¹⁰ For the first period, purchase probabilities are given by a special case of equation (5),

$$\Pr(d = j|(o, 0), 1) = \frac{\exp \left(\delta_{j1}^f - \gamma_{j1} - \phi_{j1} + \beta \mathbb{E}V((j, 1), 2) \right)}{\exp \left(\beta \mathbb{E}V((o, 1), 1) \right) + \sum_{k=1}^{J_1} \exp \left(\delta_{k1}^f - \gamma_{k1} - \phi_{k1} + \beta \mathbb{E}V((k, 1), 2) \right)}.$$

These probabilities define $s_{ij1}, \forall i, j$ and are stored in matrix $H(1)$.

For each of $t > 1$, we compute purchase probabilities in equation (5) and calculate type-specific shares using,

$$s_{ijt} = \sum_{e \in \bigcup_{\tau=1}^{t-1} \mathcal{J}_\tau \times \{\tau\}} \Pr_i(d = j|e, t) h_e^i(t-1) + \sum_{\tau=1}^{t-1} h_{(j, \tau)}^i(t-1) \left(1 - \sum_{k \in \mathcal{J}_t} \Pr_i(d = k|(j, \tau), t) \right), \quad (6)$$

where the first term corresponds to the ‘‘active’’ purchase decision and the second term corresponds to $d_i = \emptyset$, i.e., the decision to keep holdings from the previous periods. Finally, we update the holdings matrix $H(t)$ by using the first term in equation (6) for the current-period column, $h_{(j,t)}^i(t)$, i.e.,

$$h_{(j,t)}^i(t) = \sum_{e \in \bigcup_{\tau=1}^{t-1} \mathcal{J}_\tau \times \{\tau\}} \Pr_i(d = j|e, t) h_e^i(t-1), \quad (7)$$

and

$$h_{(j,\tau)}^i(t) = h_{(j,\tau)}^i(t-1) \left(1 - \sum_{k \in \mathcal{J}_t} \Pr_i(d = k|(j, \tau), t) \right) \quad (8)$$

for all $\tau < t$.

Finally, the aggregate market share of product j is computed as the weighted sum of product shares for all consumer types:

$$s_{jt} = \sum_{i=1}^N w_{it} s_{ijt}, \quad (9)$$

¹⁰We discuss the initial conditions problem in Section 6.

where w_{it} is the weight of consumer type i in the population.

4.3 Mean flow utility and estimation algorithm

We assume that all consumer types have the same preferences for product characteristics (x_{jt}, ξ_{jt}) ; i.e., consumer types differ only with respect to price sensitivity. Our model predicts aggregate market shares as a function of $((\delta_{k\tau}^f)_{k \in \mathcal{J}_\tau})_{0 \leq \tau \leq T}$ for any given vector of $\bar{\alpha}_p = (\alpha_{1p}, \dots, \alpha_{Np})$. Let $\hat{s}_{jt}(((\delta_{k\tau}^f)_{k \in \mathcal{J}_\tau})_{0 \leq \tau \leq T}, \bar{\alpha}_p)$ be the predicted market share.

Our first step is to recover δ_{jt}^f for all products in all time periods. Since we don't observe consumer flow utilities in the data, we begin with arbitrary starting values for δ_{jt} and estimate $(\hat{\gamma}_{0,g}, \hat{\gamma}_{1,g})$. We also estimate two more pairs of parameters—one for the evolution of handset prices and another one for the evolution of service fees. Then, we fill in the extended data set for all periods and products beyond the observed terminal period T and solve the consumer dynamic programming problem by backwards induction. The solution to the consumer problem is then used to predict market shares. Mean flow utilities are then recovered by solving the following system of equations in $((\delta_{k\tau}^f)_{k \in \mathcal{J}_\tau})_{0 \leq \tau \leq T}$,

$$s_{jt} = \hat{s}_{jt} \left(\left(\left(\delta_{k\tau}^f \right)_{k \in \mathcal{J}_\tau} \right)_{0 \leq \tau \leq T}, \bar{\alpha}_p \right), \quad \forall t \in \{0, \dots, T\}, \forall j \in \mathcal{J}_t. \quad (10)$$

In particular, we employ an inversion algorithm similar to Berry et al. (1995) to update the mean flow utilities as follows:

$$\delta_{jt}^{f(1)} = \delta_{jt}^{f(0)} + \log(s_{jt}) - \log \left(\hat{s}_{jt} \left(\left(\left(\delta_{k\tau}^f \right)_{k \in \mathcal{J}_\tau} \right)_{0 \leq \tau \leq T}, \bar{\alpha}_p \right) \right),$$

where $\delta_{jt}^{f(0)}$ and $\delta_{jt}^{f(1)}$ are the current and the next iteration values of the mean flow utility of product j at time t . Using new sequences of the payoff-relevant variables we update our solution to the consumer maximization problem. Iterations repeat until value functions from two consecutive iterations are close enough.

To estimate structural parameters we use a nested fixed-point algorithm. In the inner loop, for any given sequence of the payoff-relevant variables, such as handset prices, service fees, and mean flow utilities, we numerically solve the consumer dynamic programming problem (2). The solution is then used to update aggregate market share predictions defined by equations (6) and (9). A new sequence of the mean flow utilities is then recovered by solving equation (10). To approximate continuation values in the dynamic programming problem beyond the final period observed in the data we estimate AR(1) regressions of current values of the handset price, service fee, and mean flow utility on their one-period-lagged values. Parameter estimates from the regressions are then used to predict future values of these variables for 60 periods in the future.¹¹

After convergence, we can use the recovered mean flow utilities to net out unobserved product characteristics and form moment conditions for estimation. In particular, we assume that

$$\delta_{jt}^f = \delta_{hct}^f = \alpha_0 + \alpha_h + \alpha_{ct} + \xi_{jt}, \quad (11)$$

where α_h is a vector of handset fixed effects (FE), α_{ct} is a vector of carrier-time FE, and ξ_{jt} are product-specific innovations satisfying the following assumption:

$$\mathbb{E}[\xi_{jt} | Z_{jt}] = 0, \quad (12)$$

¹¹Given the definition of time period in our model, 60 periods are equivalent to 30 years.

where Z_{jt} is a vector of instrumental variables discussed in Section 5.

After calculations in the inner loop are completed, we concentrate out linear parameters defined in equation (11) and interact product-level unobservables ξ_{jt} with instrumental variables to evaluate the generalized method of moments (GMM) criterion function. The outer loop searches for nonlinear parameter values by employing the Nelder-Mead simplex optimization method.

The main objective of our estimation algorithm is to recover structural parameters α_{ip} , α_0 , α_h , and α_{ct} , which can then be used to simulate counterfactual scenarios where carriers do not charge early termination fees.

In the heterogeneous consumer version of the model, we facilitate identification of the demographic type-specific price coefficients by including additional (micro) moment conditions based on the difference between model prediction and observed type-specific purchase probabilities. In particular, consumer i 's purchase probability for product j at time t is given by s_{ijt} in equation (6). Let \hat{s}_{ijt} denote the purchase probability for consumer type i predicted by the model and let s_{ijt} denote the corresponding market share observed in the data. Then, we define additional moment conditions using error terms

$$\nu_{ijt} = s_{ijt} - \hat{s}_{ijt}, \text{ s.t. } \mathbb{E}[\nu_{ijt}|\mathcal{I}_t] = 0; \quad (13)$$

that is, the conditional expectation of the errors in the type-specific market shares is zero given the information set at time t . In this paper we use only one instrumental variable for the micro-moment conditions, which is a constant term. To construct the GMM objective function we stack moments based on ξ and ν horizontally and employ a block-diagonal weighting matrix. The weighting matrix for the first stage GMM is calculated as $(Z'Z)^{-1}$, while the second stage optimal weighting matrix is based on the inverse covariance matrix of the individual moment conditions.

5 Instruments and identification

It is conceivable that wireless service providers observe ξ_{jt} (at least partially) prior to choosing their service fees, ϕ_{jt} , and handset prices, γ_{jt} . We are less concerned with the early termination fees as they are typically identical for a wide range of products (in our data they differ only for smartphones vs non-smartphones) and change very infrequently.

In order to address the endogeneity problem we construct several instrumental variables similar to Berry et al. (1995). We don't observe any provider-specific characteristics except for the identity of the carrier. By using carrier-time effects we control for the nation-wide quality of each carrier's service at any given time period. For handsets, we observe very detailed information as discussed in Section 3. Our instrumental variables measure the intensity of competition facing each of the products as the number and average characteristics of similar handsets offered by competitors. For a given product, a larger number of substitutes as well as their closer proximity in the characteristics space should negatively affect price-cost margins.

Similar handsets are defined based on whether a given handset is a smartphone, availability of camera, type of OS vendor, handset form factor, and the total number of observable features (e.g., GPS capabilities, radio, Bluetooth, Java, built-in storage, wifi data access, etc.). We construct four distinct instrumental variables. The first variable is the number of identical handsets as defined by their characteristics currently offered by the competing carriers. The second variable is the average (across rival products) age

of similar handsets as measured by the number of months since the introduction of the product. The third variable is the average self-reported consumer satisfaction by similar rival products. The fourth instrumental variable measures the total number of handsets brought to the market by the same original equipment manufacturer (OEM) in a given time period.

Identification of parameters in our model is based on several assumptions. First, we assume that all consumers discount the future at the same rate of $\beta = 0.95$. Second, our specification of mean population utility from a handset-carrier combination assumes that all consumer types have identical preferences for the attributes of handsets and carriers. In other words, we do not allow for random coefficients on handset and carrier-time dummies. The identification of each consumer type’s price sensitivity relies on micro-moments.

One important question is the uniqueness of the mapping between observed market shares and mean flow utilities. Similar to Gowrisankaran and Rysman (2012), we allow for repeated consumer purchases over time. Therefore, while the information structure of our problem is similar to Berry et al. (1995), consumers may choose to buy multiple products over time. This prevents us from using the uniqueness proof as in Berry (1994) because it would require substitution between all products. As a result, we proceed by assuming uniqueness of the vector of mean flow utilities that makes observed market shares equal to the model predictions. Extensive testing reveals that for all trial parameter values and all initial starting values for the mean-flow utility vector, the algorithm always converges to the same solution.

Since we observe data from the first quarter of 2005 until the third quarter of 2012 there are obvious initial conditions and terminal period problems. As discussed in Section 4, our main specification assumes that after the terminal period product characteristics evolve according to a deterministic process (3). In Section 6, we also discuss parameter estimates under alternative assumptions that products and their characteristics stay constant after 2012 or evolve stochastically. We also conduct similar robustness checks for the initial conditions problem.

6 Estimation results

We begin by presenting results from several static discrete choice specifications. The main purpose of these regressions is to illustrate the effect of instrumental variables on the parameter estimates. Next, we estimate a representative consumer and several heterogeneous consumer versions of our structural model. Discussion of the estimation results follows.

Static model results. In Table 2, we report estimation results for several specifications of a simple static model. In this model, every period consumers decide whether to purchase one of the products or to choose an outside option. Every time a product is purchased, a consumer has to pay a handset price and a service fee. Early termination fees were not used in estimation. The mean utility from each handset-carrier combination was obtained following Berry (1994). Specifications (1), (3), and (5) report OLS estimates. Results from the instrumental variable regressions are presented by specifications (2), (4), and (6). Results reported in (1) and (2) allow for difference in coefficients on the handset price and service fee variables (unrestricted), while the estimates listed in columns (3) through (6) restrict coefficients on both monetary variables to be the same (restricted).

We control for product and carrier-time fixed effects or handset and carrier-time fixed effects. The bottom part of Table 2 reports F-statistics from the first-stage regressions in the IV specifications. Additional static specification results are reported in Appendix B (Tables 16 and 17).

Table 2 suggests the presence of the endogeneity problem for handset price, service fee, and total cost variables. OLS specifications estimate the positive price coefficients on the service fee and total cost variables. Endogenous variables in the IV specifications are instrumented with (1) total number, (2) average age, and (3) average consumer satisfaction by products of competing carriers, as well as with (4) the total number of handsets produced by the same OEM. Price coefficients in the IV regressions are estimated to be negative as expected. Given the endogeneity concern, we use similar instrumental variables to form moment conditions in estimation of the structural model.

Table 2: Results from static model specifications: OLS vs IV, 16,408 observations

parameters	unrestricted		restricted			
	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV
handset price, P_{jt}	-4.591	-15.565				
(s.e.)	(0.262)	(14.890)				
service fee, p_{jt}	1.021	-17.273				
(s.e.)	(0.091)	(4.345)				
total cost, $p_{jt} + P_{jt}$			0.340	-17.190	0.755	-9.574
(s.e.)			(0.085)	(4.296)	(0.094)	(2.846)
constant	-8.070	1.702	-8.148	1.787	-3.640	1.368
(s.e.)	(0.046)	(2.216)	(0.046)	(2.092)	(0.173)	(1.396)
product fixed effect	yes	yes	yes	yes	no	no
carrier-time fixed effect	yes	yes	yes	yes	yes	yes
handset fixed effect	no	no	no	no	yes	yes
first-stage statistics						
F statistic, P_{jt}		11.76				
(p-value)		(0.000)				
F statistic, p_{jt}		14.34				
(p-value)		(0.000)				
F statistic, $P_{jt} + p_{jt}$				18.61		6.58
(p-value)				(0.000)		(0.000)

Notes: The total cost variable is the sum of handset price and service fee; instrumental variables include average number, average age, and average consumer satisfaction for products offered by the rival carriers as well as total number of handsets produced by the same OEM over time.


Structural model results. All reported results from the structural model are second-stage optimal GMM parameter estimates. First-stage results are available upon request. Table 3 summarizes estimation results from a representative consumer version of the structural model (in column 2). For convenience of comparison in column (1) we also reproduced estimation results from a similar static model. The bottom part of the table reports mean, median, and standard deviation for the aggregate price elasticity of demand. The last row presents results from the appropriate version of the overidentifying restrictions test. Neither the static nor dynamic model can be rejected at a reasonable significance level.

Table 3: Second-stage optimal GMM parameter estimates and elasticity predictions

parameter	parameter estimates	
	(1) static	(2) dynamic
price coefficient, α_p	-9.574	-8.163
(s.e.)	(2.846)	(3.014)
carrier-time fixed effects	yes	yes
handset fixed effects	yes	yes
<i>service fee elasticity</i>		
average	-3.477	-2.967
median	-3.558	-3.033
standard deviation	0.981	0.838
<i>handset price elasticity</i>		
average	-0.618	-0.524
median	-0.484	-0.410
standard deviation	0.478	0.404
Sargan stat/Hansen's J-stat	2.475	0.841
(p-value)	(0.480)	(0.359)

Estimation results from the dynamic model suggest an average own service fee elasticity of -2.97, while the static version estimate for this parameter is -3.48. Elasticity estimates with respect to handset price are in the inelastic range and constitute -0.52 and -0.62 for dynamic and static models, respectively. Such a low elasticity measures are consistent with subsidized handset prices. Figure 6 reports histograms for service fee and handset own price elasticity from the dynamic model. Typical estimates of the own price elasticity in the existing literature are slightly lower. For example, Ingraham and Sidak (2004) find it to be about -1.29, while Caves (2011) reports -2.1.

Figure 6: Distribution of own price elasticity for all products



Estimates from our structural model can be used to obtain predictions for consumer switching rates. We also compute implied profits from the ETF payments made by consumers who switch before their contract expires. Table 4 reports summary statistics for the churn rates (probability that a consumer switches away) and expected profits from the ETF payments by carrier.

Table 4: Monthly churn rates and revenues from ETFs by carrier, structural model

carrier	churn rates			ETF-revenues/subscriber		
	mean	median	st.dev.	mean	median	st.dev.
ATT	0.03	0.03	0.00	2.53	2.52	0.16
OTH	0.01	0.01	0.00	0.84	0.85	0.12
SPR	0.02	0.02	0.00	1.91	1.93	0.50
TMO	0.01	0.01	0.00	1.27	1.35	0.15
VER	0.03	0.04	0.01	3.14	3.13	0.08
average	0.02	0.02	0.01	1.94	1.93	0.87

Notes: Revenues from ETFs are provided at the monthly level assuming a market size of one.

On average churn rates are estimated to be about 2 percent per month, which is consistent with the anecdotal evidence often reported in economic literature on the wireless industry and media. The highest churn rates are estimated for AT&T and Verizon followed by Sprint, T-Mobile and other independent operators. Every time a consumer switches service provider before the contract expiration date an early termination fee must be paid. Estimates from our structural model suggest that wireless operators expect to earn about 1.94 dollars per month per subscriber. Similar to the case of churn rates, Verizon and AT&T earn the most, with operators classified as other earning the least. We will use estimates of profits generated by the ETF payments in Section 7 when we present results from our full equilibrium counterfactual simulations.

Initial conditions and terminal period beliefs. We conducted robustness checks for our assumptions about the initial conditions and terminal period consumer beliefs about the industry evolution. Recall that our main specification assumes every consumer holds an outside option at $t = 0$ in our sample, while the evolution of handset prices, service fees and mean flow utilities after the final period in the sample $t > T$ is perceived by consumers as an AR(1) process with zero innovations. To evaluate the effect of these assumptions on the parameter estimates we re-estimated our representative consumer specification under alternative scenarios. In all of them, we assume the wireless market emerged 15 time periods before our sample and consumers consider 60 periods forward after the last period in our sample. We experiment with an AR(1) process for the payoff-relevant variables with and without random innovations in the pre- and post-sample periods. We also re-estimate the model with constant pre- and post-sample period values of handset prices, service fees, and flow utilities. Various combinations of deterministic and stochastic evolution for these variables are tested. Overall, we find very similar estimates for the models, where the initial conditions and post-terminal period values of the payoff-relevant variables are approximated with their first and last values in the sample, respectively. For example, the highest and the lowest estimates across our experiments are -8.100 (2.850) and -8.208 (2.974), respectively.¹² We also find that assumptions about the post-terminal period expectations have a relatively stronger effect on the parameter estimates than the assumptions made about consumer holdings prior to our sample period.

Alternative sets of instrumental variables. By using a representative consumer version of the structural model we conduct extensive robustness checks for various subsets of instrumental variables. Alternative estimation results and counterfactual simulations

¹²Full estimation results are available upon request.

can be found in Appendix B in Tables 18, 19, 20, 21, and 22 and Figures 11 and 12. Results from a similar robustness exercise for static models are listed in Tables 16 and 17 in the same appendix.

Alternative specifications of instrumental variables cannot be discriminated according to the overidentifying restrictions tests because none of them is rejected. Furthermore, the t-test does not reject similarity in the parameter estimates across specifications with alternative sets of IVs. In particular, when comparing results in Table 18 in Appendix B, the t-statistic for the difference between (1) and (2) has a p-value of 0.69; when comparing (1) and (3) the p-value is 0.42; and the p-value for the difference between (2) and (3) is 0.65.¹³ Therefore, to select our main specification we used a different criterion, namely the share of products with negative predictions for marginal costs. It turns out that the price parameter estimate of -8.163 produces the smallest share of negative marginal cost predictions among all tested specifications (see Table 10 in Section 7 for further details). We use this selection criterion because estimation of the dynamic demand model in our case does not rely on any supply-side moment conditions that, if included, wouldn't admit negative cost estimates.

Next we move to heterogeneous consumer versions of the structural model. As we discuss in Section 4, to facilitate identification of the type-specific price coefficients we augment our GMM criterion function by micro-moments based on the differences between the type-specific market share predicted by the model and the corresponding purchase probabilities observed in the data, as defined by equation (13). We find that including additional micro-moments significantly improves the precision of the estimates.

Table 5: Estimation results for heterogeneous consumers


type	(1)	(2)	(3)	
	income	age	age < 45	age ≥ 45
income, < 50K	-8.777		-10.483	-12.953
(s.e.)	(2.103)		(2.246)	(2.239)
income, ≥ 50K	-7.805		-5.283	-11.152
(s.e.)	(2.102)		(2.043)	(2.248)
age, < 45		-7.068		
(s.e.)		(2.151)		
age, ≥ 45		-9.759		
(s.e.)		(2.165)		
carrier-time dummy	yes	yes	yes	
handset dummy	yes	yes	yes	
Hansen J-stat	1.419	1.357	2.526	
(p-value)	(0.492)	(0.507)	(0.283)	

Table 5 summarizes estimation results for several consumer types defined as belonging to age, income, or age-income groups. Parameter estimates for the two-type model with high- and low-income consumers are reported in column (1). Results from the two-type model with young and old consumers can be found in column (2). Column (3) summarizes estimates from our most rich four-type model with consumers tabulated into high and low types according to their age and income.

Parameter estimates from the model with four consumer types suggest a slightly

¹³The t-test is computed assuming zero covariance between the coefficients.

higher own elasticity with respect to service fee and handset price than a representative consumer version does. Figure 7 overlays elasticity histograms for one- and four-type models. Differences in the elasticity with respect to service fees appear slightly larger.



We also compare estimates for carrier-time dummy variables for one-type and four-type models. Both models suggest a constant decline in the average quality of each carrier over time. This is consistent with the improvements in the outside option represented by any non-contract-based telecommunication service such as Voice over Internet Protocol and various internet messengers. Table 6 provides summary statistics for the estimates. Overall, we find carrier-time effects in the two models are closely correlated with a correlation coefficient of 0.997. The heterogeneous-type model estimates slightly lower parameter values.

Table 6: Summary statistics for carrier-time effects for one- and four-type models

model	mean	median	min	max	s.d.
one-type	-2.46	-2.38	-5.28	0.00	1.23
four-type	-2.40	-2.32	-5.19	0.00	1.21

Recall that in our theory model, quality levels of the competing firms are perfectly negatively correlated. As a result, the average over time quality measure appears to be very similar for both firms. Relative homogeneity of the products (streams of utility flows) in a dynamic model, in turn, does not allow firms to set high margins. Elimination of ETFs allows consumers to choose different products over time because each of them is sufficiently differentiated at a given point. The difference in product differentiation determines the potential price response in the counterfactual scenario without long-term contracts.

Table 7: Correlations in carrier-time effects for one- and four-type models

	one-type model					four-type model				
	ATT	OTH	SPR	TMO	VER	ATT	OTH	SPR	TMO	VER
ATT	1.00					1.00				
OTH	0.93	1.00				0.92	1.00			
SPR	0.98	0.97	1.00			0.98	0.97	1.00		
TMO	0.99	0.93	0.97	1.00		0.99	0.91	0.96	1.00	
VER	0.99	0.95	0.99	0.99	1.00	0.99	0.94	0.99	0.98	1.00

Table 7 reports correlations in estimated carrier-time dummies. Estimates of the average product quality appear to be closely correlated across carriers. If the difference in product differentiation with and without ETF is what affects firms' markups and prices, we do not expect to see a dramatic price change after the long-term contracts are eliminated.

7 Counterfactual simulations

We begin by describing several partial equilibrium counterfactual simulations. We refer to these simulations as partial because they rely on the demand-side parameter estimates only. In these experiments, we do not allow service providers to re-optimize their service fees after the counterfactual change. Instead, we calculate the compensating proportional change in the service fees that would offset gains in consumer welfare after the elimination of early termination fees.

7.1 Partial equilibrium counterfactuals

Our partial equilibrium analysis assumes the following four counterfactual scenarios. First, we set all ETFs to zero and solve the consumer's dynamic problem holding handset prices and service fees fixed (line "No, purchased at obs. prices"). Second, we assume that when ETFs are eliminated, handset prices are no longer subsidized, i.e., consumers have to pay the full price when buying a new phone (line "No, purchased at new prices"). Third, for each handset, we calculate a hypothetical rental per-period price. To do this, we estimate the per-period depreciation rate using prices reported in the survey. Next, we remove two sources of dynamics by setting ETFs to zero and allowing consumers to rent a handset on a per-period basis (line "No, rented"). Here, we assume that a handset can be rented every period at a price equal to the value of its depreciation. Finally, we allow consumers to rent a handset and firms to set early termination fees. As a result, every period consumers can choose among any handset offered by their service provider. However, to avoid paying ETFs, when switching to a different service provider, a consumer must wait until contract expiration (line "Yes, rented"). Summary statistics for the changes in the consumer welfare and market shares are reported in Table 8.

Table 8: Summary statistics for changes in consumer welfare (value functions for each holding) and market shares relative to the observed outcomes, one-type model

counterfactual scenario		mean	p50	min	max	sd
ETFs	handset	change in value functions				
No	purchased at obs. prices	0.76	0.73	0.67	0.98	0.05
No	purchased at new prices	0.48	0.47	0.38	0.68	0.04
No	rented	1.16	1.13	1.02	1.43	0.06
Yes	rented	0.19	0.19	0.16	0.21	0.01
ETFs	handset	change in market shares				
No	purchased at obs. prices	0.48	0.45	-0.76	2.25	0.35
No	purchased at new prices	0.63	0.63	-0.89	3.97	0.59
No	rented	0.70	0.41	-0.87	18.50	1.08
Yes	rented	0.31	0.12	-0.60	8.01	0.68

Notes: Service fees are fixed at observed level; changes are calculated using $\Delta(v)_i = (v(i) - V(0))/V(0)$, where i denotes the scenario, $v(i)$ is the variable of interest and $V(0)$ is the factual outcome.

As expected, the highest increase in consumer welfare (116 percent) can be achieved by removing ETFs and simultaneously allowing consumers to rent or lease their handset at its depreciation cost. Interestingly, the rental option alone, which eliminates dynamics generated by durability of the handsets, can improve consumer welfare by 19 percent. Elimination of the ETFs under current handset prices can increase consumer welfare on average by 76 percent. Since it is not very realistic to expect the same subsidized handset prices when ETFs are abandoned, we calculate the expected increase in consumer welfare under the scenario of no subsidized handsets. The resulting welfare improvement is less impressive and is equal to 48 percent. Naturally, the increase in consumer utility from the available products would increase market shares with the change in product-specific market shares ranging between 31 and 70 percent on average.

The welfare comparison table for the four-type heterogeneous consumer model can be found in Appendix B in Table 21. The heterogeneous consumer version of the model, due to its slightly higher estimated aggregate price sensitivity, suggests higher consumer gains in each of the scenarios. For example, elimination of ETFs at observed handset prices improves consumer welfare by 90 percent instead of 76 percent in the representative consumer model. At new handset prices, the multi-type model suggests the difference of 58 percent improvement relative to 48 percent in the representative consumer case. Similarly, the pure rental option can increase consumer welfare by 22 percent instead of 19 percent.

To illustrate the relationship between demand elasticity and consumer welfare predictions we also repeat the same exercise for alternative versions of the representative consumer model summarized in Appendix B in Table 18. The corresponding welfare comparisons are reported in the same appendix in Tables 19 and 20. Lower estimated price elasticity (see Figures 11 and 12) results in more modest welfare improvements. For example, the best-case scenario, when consumers do not face ETFs and can rent their handsets, can improve consumer welfare between 66 and 84 percent. In a more realistic scenario, when ETFs are eliminated and handset prices are no longer subsidized, consumer welfare increases by 27 to 34 percent.

ETFs versus service fees. Before we specify the supply side for our model and make additional restrictive assumptions, let us provide a quick analysis of the compensating service fee change, which would offset welfare gains from the elimination of ETFs. In particular, suppose all service fees increase proportionally, i.e., an increase by x percent to $p_{jt}(x)$ is defined by $p_{jt}(x) = (1 + x)p_{jt}$ for all j, t . Then, we can solve for x^* such that the average difference between the old and new (no ETFs but higher fees) value functions is zero. The value x^* would then determine the price change that makes consumers indifferent between the original situation and the new one. Table 9 reports results of the simulation for our representative and heterogeneous four-type consumer models.

Table 9: Change in service fees offsetting consumer gains from ETF elimination, %

type of compensating change	one-type model	four-type model
increase in service fees at obs. h-set prices	42.59	41.21
increase in service fees at new h-set prices	31.70	29.60

Notes: Offsetting price increase is computed such that the difference between consumer value functions before the ETF elimination and consumer value functions after the ETF elimination with corresponding proportional change in service fees is zero on average.

While the heterogeneous consumer model predicts slightly higher elasticity with respect to service fees, the compensating increase in these fees is slightly lower. To investigate the relationship between elasticity estimates and the offsetting price increase further, we conduct simulations for alternative specifications of the representative consumer model reported in Appendix B, Table 18. The results are summarized in Table 22. The smaller estimated price sensitivity is associated with a weakly larger offsetting service fee increase. However, the differences do not appear significant.

7.2 Supply of wireless services

To recover the cost structure of the wireless service providers we make a major simplifying assumption about producer rationality. We consider a version of the supply-side model that assumes forward-looking firms are capable of predicting future sequences of marginal costs and accounting for dynamic consumer behavior. At the same time, we impose a restriction on the firms' behavior by focusing on the open-loop strategies. While this model has clear limitations in a quickly developing market like the wireless industry, we choose the commitment case because of its computational tractability.

Notation. Time is discrete; the horizon is infinite: $t = 0, \dots, T, \dots, \infty$, where T denotes the terminal period in our sample. Let \mathcal{F} be the set of firms. Due to the perfect foresight assumption on the demand side, the market share of every product depends on the entire vector of characteristics of all products in all time periods. Let \mathcal{J}_t denote the total number of products available for purchase at time t and define

$$\mathbf{P} = \left((p_{jt})_{j \in \mathcal{J}_t} \right)_{t=0}^{\infty}.$$

We denote products offered by firm f at time t by $\mathcal{J}_{f,t}$. Since we keep product attributes other than service fee fixed throughout our simulations we denote the period t market share of product j purchased at time t' with $s_{t,(j,t')}(\mathbf{P})$, i.e., as a function of the price vector only. This share is given by the fraction of consumers who purchased a particular

handset-carrier combination (j, t') and have not replaced it yet with a new product by the end of t . Note that the current period aggregate market share of product j consists of consumers who bought it in any of the time periods since it became available on the market. Due to sticky contract prices, to calculate profits of the wireless carriers we need to know the shares of consumers for each potential holding state. This is because consumers who purchased a product at t' pay service fee $p_{jt'}$ and not the going market price p_{jt} . Therefore, a wireless carrier profit per subscriber is given by

$$\pi_f = \sum_{t=0}^{\infty} \beta^t \sum_{j \in \mathcal{J}_{f,t}} \sum_{t'=0}^t s_{t,(j,t')}(\mathbf{P})(p_{jt'} - c_{jt'}). \quad (14)$$

The open-loop case with forward-looking firms and no sticky prices. Note that the total share of product j at time t is $s_{jt} = \sum_{t'=0}^t s_{t,(j,t')}(\mathbf{P})$, and it combines all purchase decisions made from the initial product introduction period until t . As a result, first-order conditions of the profit function (14) with respect to a product price choice in a given time period would involve differentiation of all previous holding probabilities $s_{t,(j,t')}$. To make the model computationally tractable we make an additional simplifying assumption that firms' profits are derived in a non-sticky price environment. In other words, all consumers holding product (j, t') at time t pay the going market price for this product p_{jt} and not their original contract price $p_{jt'}$. Under this assumption, we can rewrite equation (14) as follows:

$$\pi_f = \sum_{t=0}^{\infty} \beta^t \sum_{j \in \mathcal{J}_{f,t}} s_{jt}(\mathbf{P})(p_{jt} - c_{jt}). \quad (15)$$

Further, since in our demand model we assume that consumers believe in the AR(1) evolution of the payoff-relevant variables after the final period observed in the data, we do not allow firms to choose prices beyond T . Consumers then perceive the future values of service fees after T as following a new AR(1) process induced by a new price vector at $t \leq T$. Then, the profit maximization problem of firm $f \in \mathcal{F}$ can be written

$$\max_{\mathbf{P}_f} \sum_{t=0}^T \beta^t \sum_{j \in \mathcal{J}_{f,t}} s_{jt}(\mathbf{P})(p_{jt} - c_{jt}), \text{ where } \mathbf{P}_f = \left((p_{jt})_{j \in \mathcal{J}_{f,t}} \right)_{t=0}^T. \quad (16)$$

Timing of the game is as follows. Before the game starts, at period $t = -1$, all firms simultaneously set all their prices for $t = 0, \dots, T$. These prices are then observed by the consumers. Consumers predict future prices for periods $t > T$ by using a simple AR(1) process based on the current price sequences in the sample. Each firm maximizes the profit from all of its products, and all operators classified as "Other" jointly maximize their total profit.

First-order conditions for product (j, t) are then

$$\sum_{t=0}^T \beta^t \left(\sum_{k \in \mathcal{J}_{f,t}} \frac{\partial s_{kt}(\mathbf{P})}{\partial p_{jt}} (p_{kt} - c_{kt}) + s_{jt}(\mathbf{P}) \right) = 0. \quad (17)$$

Let $D_f(\beta)$ and $S_f(\beta)$ denote the properly discounted matrix of market share derivatives and a vector of market shares for all products offered by carrier f , respectively. Let \mathbf{c}_f

denote a vector of marginal costs for all products of firm f . Then, we can calculate \mathbf{c}_f by solving a system of linear equations defined by (17), i.e.,

$$\mathbf{c}_f = D_f^{-1}(\beta) \left(D_f(\beta) \mathbf{P}_f + S_f(\beta) \right).$$

This system has the number of equations equal to the total number of unknowns, and both are equal to the total number of products brought to the market by carrier f over the period in the sample.

We solve a system of first-order conditions (17) for product-level marginal costs. Table 10 reports statistics on the proportion of positive marginal cost predictions for our main representative consumer specification and two alternative specifications for the one-type model. Estimation results for these two alternative specifications are reported in Table 18 in Appendix B, where our main model is listed as specification (3).


Table 10: Shares of positive marginal cost estimates by specification by carrier

carrier	N	share of positive marginal costs, $mc_{jt} > 0$		
		main, $\alpha_p = -8.16$	spec. (2), $\alpha_p = -6.47$	spec. (1), $\alpha_p = -5.43$
ATT	4,046	0.99	0.97	0.96
OTH	3,534	0.93	0.86	0.82
SPR	2,668	0.97	0.94	0.92
TMO	2,524	0.99	0.97	0.96
VER	3,873	0.99	0.98	0.96

It is worth mentioning our selection criteria for the main specification and for those reported primarily in the Appendix. By conducting robustness checks for alternative sets of instrumental variables, we estimate several specifications of our structural model. As we already mentioned in Section 6 above, neither the overidentifying restrictions test nor the t-test for difference in estimated coefficients can be used to discriminate between these alternatives. However, as our model does not have any supply side, we used the proportion of negative marginal cost predictions as a criterion to choose our main specification. That being said, we replicate most of the counterfactual simulations for all specifications and report the results in Appendix B.

The distribution of estimated marginal costs for all products in all periods as well as the evolution of average marginal costs and price cost margins over time are presented in Figure 8.


Figure 8: Marginal costs and price-cost margins



Note: Averaged across products, negative values of marginal costs excluded.

Figure 9 reports the marginal cost and price-cost margin evolution for all carriers and all time periods. Marginal costs estimated for the group of carriers labeled as “Other” appear to be the lowest and their price-cost margins consequently the highest.

Figure 9: Marginal costs and price-cost margins by carrier by time



Note: Averaged across products, negative values of marginal costs excluded.

Our intuition behind the finding that smaller carriers have lower marginal costs is that most of the smaller carriers originated as subsidiaries of the big four. They often rely on the infrastructure (cell towers) built by their bigger rivals. It is possible that lower marginal costs for these carriers reflect not only their lower quality of service but also lower infrastructure maintenance costs.

7.3 Full-equilibrium counterfactual simulations

To find the optimal level of service fees under the no-ETF scenario, we discretize monthly prices such that the minimum change in monthly fee constitutes one dollar. Then, we perform a search for the profit-maximizing price over the set of products that belongs to the same service provider. The algorithm circles through each product in the data and searches for the level of service fee that would maximize profit of the wireless carrier that offers this product. Termination of the algorithm occurs if in a given loop (over all products) no more improvement can be found. While we aggregate several smaller carriers

into a separate group “Other”, in simulation we assume that this group maximizes its joint profit, i.e., it acts as if “Other” is yet another carrier.

Table 11 reports results from the full equilibrium counterfactual simulations. After ETF elimination wireless service providers increase their service fees by 2 to 5 percent. The largest increase in prices is predicted for Verizon (5.17 percent), followed by AT&T (4.2 percent) and Sprint (3.93 percent). T-Mobil and smaller carriers increase prices the least. There is no clear pattern in the price change over time. The right panel of Table 11 reports change in profits in percentages. Note that the change corresponds to the profits derived from service fees only. Profits from ETF payments are not included (see the discussion below).

Table 11: Change in service fees and carrier profits without ETF, optimal prices

time	% change in service fees					% change in profits				
	ATT	OTH	SPR	TMO	VER	ATT	OTH	SPR	TMO	VER
2005-h1	4.53	1.87	5.22	3.77	4.29	83.16	133.14	126.26	133.04	73.12
2005-h2	4.40	1.77	4.82	3.02	5.06	71.32	114.38	107.00	116.60	61.55
2006-h1	4.65	1.95	4.90	3.11	5.07	67.32	111.90	103.05	110.55	58.24
2006-h2	4.25	2.03	4.84	3.07	4.91	59.39	102.73	94.64	99.71	50.98
2007-h1	4.02	2.05	4.60	2.94	4.58	51.23	91.47	85.15	88.74	42.88
2007-h2	3.79	1.87	3.92	2.61	4.48	45.37	82.79	78.22	80.72	37.12
2008-h1	3.77	1.84	3.57	2.75	4.18	40.30	75.16	71.90	73.70	32.66
2008-h2	3.61	1.95	3.47	2.68	4.12	37.73	71.45	68.63	70.42	30.33
2009-h1	3.50	2.03	2.87	2.62	3.96	35.79	69.75	67.33	68.85	28.57
2009-h2	3.50	2.34	3.05	2.52	5.04	34.41	68.44	66.06	67.42	36.58
2010-h1	3.50	2.05	2.42	2.19	5.84	39.61	61.61	59.38	60.51	49.47
2010-h2	5.02	1.71	2.50	1.95	6.39	55.67	53.76	51.56	52.75	53.19
2011-h1	5.75	1.85	2.64	2.06	7.51	61.59	54.15	50.53	53.04	61.54
2011-h2	3.98	2.88	4.79	3.74	5.31	33.06	78.83	113.57	128.00	38.29
2012-h1	4.83	3.28	5.41	3.39	6.81	32.51	81.19	116.14	128.35	37.42
average	4.20	2.10	3.93	2.83	5.17	49.90	83.38	83.96	88.83	46.13

According to the profits collected via service fees, our model predicts the largest beneficiaries of the ETF elimination are smaller carriers, T-Mobil and Sprint. AT&T and Verizon profits increase by 46 to 50 percent, while other carriers gain between 84 and 89 percent.

Table 12 reports consumer welfare under three alternative scenarios. The first two columns report factual consumer surplus as measured by the dynamic value function and its monetary equivalent for major wireless operators. The next three columns labeled “no ETF, old prices” illustrate consumer surplus in levels and in dollars as well as percentage increase in consumer utility as compared to the factual situation. The last three columns report consumer value functions and the ultimate change in consumer welfare after wireless service providers adjust prices in a new equilibrium.

Table 12: Changes in consumer value functions after ETF elimination

carrier	factual $\mathbb{E}[V_i]$		no ETF, old prices			no ETF, new prices		
	level	\$ value	level	\$ value	% dif.	level	\$ value	% dif.
ATT	19.18	2,349.46	33.73	4,132.79	75.90	32.12	3,934.64	67.47
OTH	19.01	2,328.91	33.55	4,110.00	76.48	31.95	3,913.67	68.05
SPR	19.24	2,357.66	33.81	4,141.61	75.67	32.18	3,942.81	67.23
TMO	19.12	2,342.14	33.67	4,125.16	76.13	32.06	3,927.57	67.69
VER	19.24	2,357.71	33.80	4,141.41	75.65	32.18	3,942.59	67.22
average	19.16	2,347.18	33.71	4,130.20	75.97	32.10	3,932.26	67.53

Our estimates suggest that the average individual consumer valuation of the wireless market is equivalent to 2,350 dollars. There is very little variation in consumer surplus derived from different carriers. Elimination of ETFs increases consumer value functions by 76 percent such that a new lifetime value of the wireless market increases to 4,130 dollars. When the service providers finally adjust their prices in a new equilibrium, consumer welfare declines but remains 67 percent higher than in the factual situation. Table 12 shows that consumers benefit in the new equilibrium without early termination fees.

Table 13: Wireless carriers' monthly profits under alternative scenarios

profit sources and comparison	ATT	OTH	SPR	TMO	VER
<u>Factual</u>					
Profits from service fees	2.55	0.68	1.54	1	3.37
Revenues from ETF payments	2.53	0.84	1.91	1.27	3.14
Total, factual	5.08	1.52	3.45	2.27	6.51
<u>No ETF, old prices</u>					
Profits from service fees	3.49	1.09	2.59	1.68	4.53
Revenues from ETF payments	0.00	0.00	0.00	0.00	0.00
Total, No ETF, old service fees	3.49	1.09	2.59	1.68	4.53
% of factual without ETF payments	136.86	160.29	168.18	168.00	134.42
% of factual with ETF payments	68.70	71.71	75.07	74.01	69.59
<u>No ETF, new prices</u>					
Profits from service fees	3.83	1.22	2.85	1.86	4.91
Revenues from ETF payments	0.00	0.00	0.00	0.00	0.00
Total	3.83	1.22	2.85	1.86	4.91
% of factual without ETF payments	150.20	179.41	185.06	186.00	145.70
% of factual with ETF payments	75.39	80.26	82.61	81.94	75.42
Cost of ETF to rationalize "No ETF" policy	1.25	0.30	0.60	0.41	1.60
% of No ETF, old prices	109.74	111.93	110.04	110.71	108.39

Notes: Profits are computed for a market of size one and expressed in USD.

Table 13 reports additional information about profits of the wireless service providers. Recall that in Table 11 we show that profits derived from monthly service fees increase by 50 to 90 percent after ETFs are eliminated. However, if we add expected ETF payments to the carrier profits the situation is different. If the cost of processing each ETF payment is zero, elimination of the long-term contracts reduces total profits of wireless carriers by 17 to 25 percent. Line "% of factual with ETF payments" in Table 13 reports profits in a new equilibrium without early termination fees and new service fees as a percentage of the factual profits from both service fees and ETF payments.

Since it is unlikely that the marginal cost of processing ETF payments is zero, we calculate the minimum marginal cost of an ETF payment such that the wireless service providers are better off in an equilibrium without ETFs. These costs are reported in line “Cost of ETF to rationalize “No ETF” policy.” For example, to receive higher profits without ETFs the monthly cost of processing such payments should exceed 1.60 dollars for Verizon, 1.25 dollars for AT&T, 60 cents for Sprint, and less than 50 cents for smaller carriers and T-Mobile. The last line of Table 13 shows that after ETFs are eliminated producers gain an additional 8 to 12 percent in profits by adjusting their service fees.

Robustness analysis. There are two remaining issues to address: the effect of sticky contract prices and the joint profit maximization assumption for the whole group of smaller carriers. Recall that to invert out product-level marginal costs we assume sticky contract prices away. This assumption has been made for computational reasons and only to back out the costs. For all counterfactual simulations we allow for sticky contract prices.

To see the effect of contractually fixed service fees we simulated another full-equilibrium counterfactual where producers behave as if all of their consumers pay going market prices, i.e., no sticky prices.¹⁴ Table 14 reports summary statistics for the difference in the new equilibrium prices under alternative assumptions on producer profit formation.

Table 14: The effect of sticky prices on the optimal service fees

assumption	mean	median	min	max	sd
Sticky prices	375.94	385.45	0.00	697.06	105.25
Going market prices	375.96	385.36	0.00	697.06	105.26
Difference (levels)	-0.02	0.00	-4.00	6.00	0.44

We find a small effect of sticky prices on the optimal choice of service fees by wireless carriers. Prices chosen when producers maximize profits in a contractually fixed service fee environment are very similar to prices chosen when producers maximize profits from all their consumers paying the current market price. In particular, we find that only about 15 percent of product prices are different and the difference is by at most six dollars. Under the assumption of non-sticky prices the optimal service fees appear to increase slightly more.

As we have discussed earlier in this section, to solve for marginal costs at the product level, we make an assumption that several small carriers classified into a group “Other” behave as if they are maximizing a joint profit from all of their products. There are several carriers classified into this group (e.g., Boost Mobile, Cricket, MetroPCS, Virgin Mobile, etc.). To see the role of the joint profit maximization assumption we simulate a full equilibrium counterfactual where all products from the group “Other” maximize individual profits. Table 15 reports the difference in the equilibrium prices and profits.

¹⁴On the demand side consumers still perceive prices they pay as fixed in the contract.

Table 15: Service fees under individual vs joint maximization for OTH

carrier	average service fees			% change in service fees		
	(1) factual	(2) joint	(3) individual	(1) vs (2)	(1) vs (3)	(2) vs (3)
ATT	375.88	390.70	390.37	4.25	4.16	-0.09
OTH	311.85	318.54	307.19	2.32	-1.58	-3.78
SPR	363.89	376.90	376.70	3.76	3.70	-0.06
TMO	380.41	390.18	390.02	2.75	2.70	-0.05
VER	384.86	404.39	403.99	5.44	5.33	-0.10
average	363.38	376.14	373.66	3.70	2.86	-0.82

Under the joint profit maximization assumption, product prices in the group “Other” increase by 2.32 percent on average, while under individual product profit maximization, the price level in this group declines by 1.58 percent relative to the factual outcome. Thus, the average price increase under individual product profit maximization is about 1 percent lower than it would be when joint profits are maximized. Importantly, the main difference occurs for the group “Other” itself, where alternative assumptions about profit maximization can result in almost a 4 percent difference in prices. Prices of the products offered by larger carriers are affected only slightly and the difference does not exceed a tenth of a percentage point. Differences in profits of the big carriers under alternative profit maximization assumptions are also small and stay below 1 percent. Therefore, we conclude that the error we incur when solving for marginal costs under the non-sticky price assumption is minor and is unlikely to affect our main findings.

8 Conclusions

We hypothesize that elimination of the long-term contracts in the US wireless industry at the end of 2015 represent a transition to a new equilibrium without ETFs. To evaluate the social welfare implications of this transition we develop and estimate a dynamic empirical model of consumer demand for wireless products defined as a handset-carrier combination.

We motivate our empirical analysis using a theoretical model where firms choose whether or not to impose switching costs on consumers in the form of ETFs and then compete in service fees. Theoretical findings suggest two potential equilibria where either all firms impose ETFs or no firm does so. We find that due to the forward-looking behavior of consumers, ETFs intensify competition between wireless service providers and impose substantial downward pressure on the equilibrium prices. Importantly, consumers and service providers may rank these equilibria differently. For example, ETFs need not be necessarily harmful to consumers. Higher switching costs can make two competing products relatively homogeneous from the consumer point of view if per-period flow utilities from the products are sufficiently negatively correlated. Increased homogeneity in turn restricts firms’ ability to charge high markups. Equilibrium prices can be so low that consumers may be better off in an equilibrium with ETFs. By the same logic, lower margins may reduce producer profits below what can be earned in an equilibrium without ETFs. Of course, both predictions rely on a stylized Hotelling’s linear city model where product qualities are perfectly negatively correlated. In reality we should not expect such a strong negative correlation in competing products’ quality.

The empirical version of our model accounts for multiple competing service providers, a finite level of ETFs, the multi-product nature of the wireless service providers, correlation in per-period flow utilities across carriers, and other important details that were omitted

in our theoretical model. Estimation results suggest a significant increase in consumer surplus as a result of the ETF elimination. To offset this increase in surplus, service fees have to increase by at least 30 to 40 percent depending on whether handsets are sold at subsidized or new prices. At the estimated parameter values, the predictions for average monthly churn rates are about 2 percent, and the wireless operators collect extra revenues from the ETF payments in the amount of 0.84 to 3.14 dollars per month per subscriber.

Full equilibrium counterfactuals predict an increase in service fees by 2.10 to 5.17 percent on average. Larger carriers such as Verizon and AT&T tend to raise their prices by more than Sprint, T-Mobile and other smaller carriers. Consumers are definitely better off in a new equilibrium with an average increase in welfare of about 68 percent. There is also a substantial increase in profits gathered by means of service fees, where larger carriers (AT&T and Verizon) gain about 50 percent, while the rest of the operators increase profits by more than 80 percent. However, if we account for profits collected from the ETF payments, the overall effect on the producer welfare is less clear. In particular, if monthly costs of processing ETF payments are sufficiently high (e.g., 1.60 dollars for Verizon, 1.25 dollars for AT&T, 60 cents for Sprint, and less than 41 cent for T-Mobile and smaller carriers), then producers are also better off in a new equilibrium without ETFs. In this case, we would conclude that elimination of ETFs is, in total, welfare enhancing.

References

- Victor Aguirregabiria and Aviv Nevo. Recent developments in empirical IO: dynamic demand and dynamic games. In Daron Acemoglu, Manuel Arellano, and Eddie Dekel, editors, *Advances in Economics and Econometrics*, volume 3, Econometric Society Monographs 2, pages 53–122. Cambridge University Press, 2013.
- Steven Berry. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25:242–262, 1994.
- Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63:841–890, July 1995.
- Stephen J. Blumberg and Julian V. Luke. Wireless substitution: Early release of estimates from the national health interview survey, July-December 2014. Technical report, U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, Division of Health Interview Statistics, National Center for Health Statistics, June 2015.
- Kevin W. Caves. Quantifying price-driven wireless substitution in telephony. *Telecommunications Policy*, 35(11):984998, December 2011.
- CTIA. CTIA’s wireless industry summary report, year-end 2014 results. Technical report, CTIA-The Wireless Association, 2015.
- Joseph Cullen and Oleksandr Shcherbakov. Measuring consumer switching costs in the wireless industry. Unpublished manuscript. Yale University and Harvard University., April 2010.
- Jean-Pierre Dube, Gunter Hitsch, and Peter Rossi. Do switching costs make markets less competitive? *Journal of Marketing Research*, 46(4):435–445, August 2009.

- Gautam Gowrisankaran and Marc Rysman. Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, 120(6):1173–1219, December 2012.
- Allan Ingraham and Gregory Sidak. Do states tax wireless services inefficiently? evidence on the price elasticity of demand. *Virginia Tax Review*, 24(2):249–261, Fall 2004.
- Paul Klemperer. Competition when consumers have switching costs: an overview with applications to industrial organization, macroeconomics, and international trade. *The Review of Economic Studies*, 62:515–539, October 1995.
- Kathleen Nosal. Estimating switching costs for Medicare Advantage plans. Unpublished manuscript. University of Arizona., June 2012.
- Pew Research Center. The smartphone difference. Technical report, Pew Research Center, April 2015. URL <http://www.pewinternet.org/2015/04/01/us-smartphone-use-in-2015/>.
- Pasquale Schiraldi. Automobile replacement: a dynamic structural approach. *RAND Journal of Economics*, 42(2):266–291, 2011. Summer.

Appendix A Theory

Appendix A.1 The model

We consider a simple model with two firms, 1 and 2, and a continuum of consumers. Time is discrete and runs from $t = 0$ to $t = \infty$.

Consumers. There is a mass 1 of consumers with unit demands. When consumer i consumes one unit of firm 1's product at time $t \geq 0$, they receive a utility flow of $\delta - Cx_{it} - p_{1t}$, where $\delta > 0$ is a constant quality parameter, $C > 0$ is the transport cost parameter, p_{1t} is the price set by firm 1 in period t , and $x_{it} \in [0, 1]$ is consumer i 's type in period t . If, instead, consumer i consumes firm 2's product, then she receives utility flow $\delta - C(1 - x_{it}) - p_{2t}$, where p_{2t} is firm 2's price in period t . Consumers' outside options are normalized to 0. Throughout this section, we assume that δ is sufficiently high, so that the market always remains covered; i.e., in equilibrium, no consumer ever goes for the outside option. Consumers are risk neutral and discount the future at rate $\beta_c \in [0, 1)$.

Let \mathcal{I} be the set of consumers. We assume that consumers' types at every $t \geq 0$ $(x_{it})_{i \in \mathcal{I}}$ are drawn iid from a uniform distribution on interval $[0, 1]$. We also assume that, for every $i \in \mathcal{I}$, the profile of consumer i 's types $(x_{it})_{t \geq 0}$ is iid drawn from a uniform distribution on interval $[0, 1]$. In other words, consumers' types are independent across consumers and over time.

Firms. Firms are symmetric and operate with a constant returns-to-scale technology. We normalize the constant unit cost to 0 without loss of generality. Firms are risk neutral and discount the future at rate $\beta_f \in [0, 1)$.

In addition to the pricing decisions they make every period, firms also need to decide whether to use early ETFs. If firm $i \in \{1, 2\}$ uses ETFs and consumer $j \in \mathcal{J}$ buys from firm i at time t and price p , then consumer j irrevocably commits to buying from firm i at the same price p in all subsequent periods. This is a very crude way of modeling ETFs, but this approach captures the fact that, when a consumer signs a contract involving ETFs, this consumer is unlikely to switch to a different firm (never switches to a different firm) for a certain number of periods (in all subsequent periods). By contrast, if firm i does not use ETFs, then buying from this firm at time t does not imply any subsequent commitment.

Timing. The game unfolds as follows. Before the dynamic game starts, at time $t = -1$, firms 1 and 2 simultaneously and non-cooperatively decide whether to use ETFs. This commitment is made once and for all at time $t = -1$. Next, the dynamic pricing game starts. The timing within period $t \geq 0$ is as follows:

1. Firms 1 and 2 simultaneously set p_{1t} and p_{2t} .
2. Consumers' types in period t are drawn and become common knowledge.¹⁵
3. Consumers who previously bought from a firm using ETFs purchase from the same firm, at the price they originally purchased at. Other consumers choose which firm to buy from, if any.

¹⁵The assumption that types are publicly observed is innocuous. It allows us to use a subgame-perfect equilibrium as our solution concept. We could assume instead that types are private information and solve for perfect Bayesian equilibria. The equilibrium outcome would be the same.

4. Payoffs are realized.

Equilibrium concept. We look for stationary subgame-perfect equilibria in pure strategies. Stationarity is defined as follows: for a given action profile in stage $t = -1$, firm i sets the same price in all subgames of the dynamic pricing game. Another way of seeing this is that we are looking for Markov-perfect equilibria with an empty state space. As usual, this restriction rules out collusive equilibria based on rewards and punishments.

Appendix A.2 The dynamic pricing game

We solve the game backward, by first solving for a stationary subgame-perfect equilibrium in each of the subgames starting at time $t = 0$, and then by moving back to stage $t = -1$. There are four subgames starting at time 0: the subgame in which no firm uses ETFs (the “no ETF subgame”), the subgame in which both firms use ETFs (the “ETF subgame”), and the asymmetric subgames in which one firm uses ETFs and the other one does not (the “mixed subgame”).

Appendix A.2.1 The no ETF subgame

Let (p_1^*, p_2^*) be a candidate for a stationary subgame-perfect equilibrium. Let $U_i(p_1, p_2, x)$ be the value of a consumer with type x when she buys from firm i , current-period prices are (p_1, p_2) and the consumer expects prices to be (p_1^*, p_2^*) in all subsequent periods. Then,

$$\begin{aligned} U_1(p_1, p_2, x) &= \delta - p_1 - Cx + \beta_c \mathbb{E}_y (\max \{U_1(p_1^*, p_2^*, y), U_2(p_1^*, p_2^*, y)\}), \\ U_2(p_1, p_2, x) &= \delta - p_2 - C(1 - x) + \beta_c \mathbb{E}_y (\max \{U_1(p_1^*, p_2^*, y), U_2(p_1^*, p_2^*, y)\}). \end{aligned}$$

It is straightforward to see that the marginal buyer has type $x = \frac{1}{2} + \frac{p_2 - p_1}{2C}$, which is the standard formula in a static Hotelling model. Under no ETF, a consumer’s choice in the current period does not affect their future utility flows. Therefore, the consumer’s problem is a static one. Next, let $V_i(p_1, p_2)$ be the value of firm i when current-period prices are (p_1, p_2) and future prices are expected to be (p_1^*, p_2^*) in all subsequent periods. Then,

$$V_i(p_1, p_2) = p_i \left(\frac{1}{2} + \frac{p_j - p_i}{2C} \right) + \beta_f V_i(p_1^*, p_2^*). \quad (18)$$

Again, without ETF, a firm’s choice in the current period has no impact on this firm’s future profits. Taking first-order conditions, we obtain firm i ’s best response to price p_j :

$$p_i = \frac{1}{2} (C + p_j).$$

It follows that $p_1^* = p_2^* = C$. Therefore, there is at most one stationary subgame-perfect equilibrium. Since the objective function in (18) is concave, it follows that $p_1^* = p_2^* = C$ is indeed a stationary subgame-perfect equilibrium.

In equilibrium, each firm earns $V_i(p_1^*, p_2^*) = \frac{1}{1 - \beta_f} \frac{C}{2}$. The expected consumer surplus is:

$$\sum_{t=0}^{\infty} \beta_c^t \left(\int_0^{1/2} (\delta - Cx - C) dx + \int_{1/2}^1 (\delta - C(1 - x) - C) dx \right) = \frac{1}{1 - \beta_c} \left(\delta - \frac{5}{4}C \right).$$

We summarize these findings in the following lemma:

Lemma 1. *The no ETF subgame has a unique stationary subgame-perfect equilibrium. The stationary equilibrium prices are $p_1^* = p_2^* = C$. Both firms earn $\frac{1}{1-\beta_f} \frac{C}{2}$. The expected consumer utility is given by $\frac{1}{1-\beta_c} \left(\delta - \frac{5}{4}C \right)$.*

Appendix A.2.2 The ETF subgame

Let (p_1^*, p_2^*) be a candidate for a stationary subgame-perfect equilibrium. Let $U_i(p_1, p_2, x)$ be the value of a consumer with type x when she buys from firm i and current-period prices are (p_1, p_2) . Then,

$$\begin{aligned} U_1(p_1, p_2, x) &= \delta - p_1 - Cx + \sum_{t=1}^{\infty} \beta_c^t \int_0^1 (\delta - Cy - p_1) dy \\ &= \frac{1}{1-\beta_c} \left(\delta - \beta_c \frac{C}{2} \right) - \frac{1}{1-\beta_c} p_1 - Cx. \end{aligned}$$

The second term in the first line reflects the fact that, if the consumer buys from firm 1 today, then they will be locked in with this firm forever. Similarly,

$$U_2(p_1, p_2, x) = \frac{1}{1-\beta_c} \left(\delta - \beta_c \frac{C}{2} \right) - \frac{1}{1-\beta_c} p_2 - C(1-x).$$

Therefore, the marginal consumer's type is given by:

$$x = \frac{1}{2} + \frac{1}{1-\beta_c} \frac{p_2 - p_1}{2C}.$$

We say that a buyer is free in the current period if they are not already locked in with a firm. Let $V_i(p_1, p_2)$ be the value of firm i per free buyer when current-period prices are (p_1, p_2) .

Then,

$$V_i(p_1, p_2) = \sum_{t=0}^{\infty} \beta_f^t p_i \left(\frac{1}{2} + \frac{1}{1-\beta_c} \frac{p_j - p_i}{2C} \right) = \frac{1}{1-\beta_f} p_i \left(\frac{1}{2} + \frac{1}{1-\beta_c} \frac{p_j - p_i}{2C} \right).$$

The pricing game is equivalent to a standard Hotelling duopoly game, with market size $\frac{1}{1-\beta_f}$ and transport cost $C(1-\beta_c)$. Therefore, $p_1^* = p_2^* = C(1-\beta_c)$, and each firm earns an equilibrium present discounted value of profits of $\frac{1}{2} C \frac{1-\beta_c}{1-\beta_f}$. The expected consumer surplus is given by:

$$\mathbb{E}_y \max(U_1(p_1^*, p_2^*, y)) = \frac{1}{1-\beta_c} \left(\delta - \beta_c \frac{C}{2} \right) - C - 2C \int_0^{\frac{1}{2}} y dx = \frac{1}{1-\beta_c} \left(\delta - \frac{5}{4}C + \frac{3}{4}\beta_c C \right).$$

The content of this section is summarized in the following lemma:

Lemma 2. *The ETF subgame has a unique stationary subgame-perfect equilibrium. The stationary equilibrium prices are $p_1^* = p_2^* = C(1-\beta_c)$. Both firms earn $\frac{1-\beta_c}{1-\beta_f} \frac{C}{2}$. Expected consumer utility is given by $\frac{1}{1-\beta_c} \left(\delta - \frac{5}{4}C + \frac{3}{4}C\beta_c \right)$.*

Comparing Lemmas 1 and 2, we see that prices are lower with ETFs than without. This comes from the fact that, if a firm slightly increases its price, a consumer that

purchases from this firm will have to pay this higher price in all subsequent periods. This amplifies the utility cost of price increases, and therefore makes demand more elastic. Another way of seeing this is that consumers perceive firms as being less differentiated under ETFs: with ETFs, the relevant differentiation parameter is $C(1 - \beta_c)$; without ETFs, it is C . This also implies that firms are worse off with ETFs. By contrast, consumers are better off with ETFs. On the one hand, ETFs intensify competition and induce lower prices. On the other hand, transport costs increase under ETFs, since consumers cannot purchase from the closest firm once they are locked in. The first effect dominates.

We cannot add up consumer surplus and producer surplus to obtain a measure of social welfare, since the firms and the consumers do not necessarily have the same discount factor. However, it is easy to see that, in period 0, social welfare is equal to $\delta - \frac{C}{4}$ with and without ETFs, and in all subsequent periods, social welfare is equal to $\delta - \frac{C}{4}$ without ETFs, and to $\delta - \frac{C}{2}$ with ETFs. Therefore, there is a sense in which ETFs degrade allocative efficiency.

So far, we have assumed that consumers are forward looking. One way of relaxing this assumption is to assume that consumers discount payoffs with discount factor β_c , but that they behave as if their discount factor were β_c^b . We say that consumers are myopic if $\beta_c^b < \beta_c$. Given this behavioral assumption, expected consumer surplus under ETF is

$$\frac{1}{1 - \beta_c} \left(\delta - \beta_c \frac{C}{2} \right) - \frac{1 - \beta_c^b}{1 - \beta_c} C - \frac{1}{4} C = \frac{1}{1 - \beta_c} \left(\delta - \frac{5}{4} C \right) + \frac{C}{1 - \beta_c} \left(\beta_c^b - \frac{1}{4} \beta_c \right),$$

whereas expected consumer surplus without ETFs is still $\frac{1}{1 - \beta_c} \left(\delta - \frac{5}{4} C \right)$. It follows that consumers suffer from the introduction of ETFs if and only if $\beta_c^b < \frac{1}{4} \beta_c$, i.e., if and only if they are myopic enough.

Appendix A.2.3 The mixed subgame

To fix ideas, suppose that firm 1 uses ETFs and firm 2 does not. Let (p_1^*, p_2^*) be a candidate for a stationary subgame-perfect equilibrium. We first analyze demand-side behavior. Let $U_i(p_1, p_2, x)$ be the value of a consumer with type x when they buy from firm i , current-period prices are (p_1, p_2) , and the consumer expects future prices to be (p_1^*, p_2^*) . Then,

$$\begin{aligned} U_1(p_1, p_2, x) &= \delta - Cx - p_1 + \frac{\beta_c}{1 - \beta_c} \left(\delta - \frac{1}{2} C - p_1 \right), \\ U_2(p_1, p_2, x) &= \delta - C(1 - x) - p_2 + \beta_c \mathbb{E}_y \max(U_1(p_1^*, p_2^*, y), U_2(p_1^*, p_2^*, y)). \end{aligned}$$

The first value function reflects the fact that, if the consumer buys from firm 1 at price p_1 in the current period, then they will have to purchase from the same firm at the same price in all subsequent periods. On the other hand, if the consumer buys from firm 2, then they will get to choose which firm to buy from in the next period. Let

$$\begin{aligned} u_1^* &= \frac{1}{1 - \beta_c} \left(\delta - \frac{1}{2} \beta_c C \right), \\ \text{and } u_2^* &= \delta + \beta_c \mathbb{E}_y \max(U_1(p_1^*, p_2^*, y), U_2(p_1^*, p_2^*, y)). \end{aligned} \tag{19}$$

The marginal consumer has an x such that $U_1(p_1, p_2, x) = U_2(p_1, p_2, x)$. Solving for x , we get:

$$x = \frac{1}{2} + \frac{u_1^* - u_2^*}{2C} + \frac{p_2 - \frac{p_1}{1-\beta_c}}{2C}. \quad (20)$$

In equilibrium, $p_1 = p_1^*$, $p_2 = p_2^*$, the marginal type is given by:

$$x^* = \frac{1}{2} + \frac{u_1^* - u_2^*}{2C} + \frac{p_2^* - \frac{p_1^*}{1-\beta_c}}{2C}, \quad (21)$$

and u_2^* satisfies

$$\begin{aligned} u_2^* &= \delta + \beta_c \left(\int_0^{x^*} \left(u_1^* - \frac{p_1^*}{1-\beta_c} - Cy \right) dy + \int_{x^*}^1 (u_2^* - p_2^* - C(1-y)) dy \right), \\ &= \delta + \beta_c \left(\left(u_1^* - \frac{p_1^*}{1-\beta_c} - \frac{1}{2}Cx^* \right) x^* + \left(u_2^* - p_2^* - \frac{1}{2}C(1-x^*) \right) (1-x^*) \right). \end{aligned} \quad (22)$$

Next, we turn our attention to the behavior of the supply side. Let $V_i(p_1, p_2)$ be the value of firm $i \in \{1, 2\}$ per free buyer when current prices are (p_1, p_2) and future prices are expected to be (p_1^*, p_2^*) . Then, using the definition of the marginal type (20),

$$\begin{aligned} V_1(p_1, p_2) &= \frac{1}{1-\beta_f} p_1 \left(\frac{1}{2} + \frac{u_1^* - u_2^*}{2C} + \frac{p_2 - \frac{p_1}{1-\beta_c}}{2C} \right) \\ &\quad + \beta_f \left(\frac{1}{2} + \frac{u_2^* - u_1^*}{2C} + \frac{\frac{p_1}{1-\beta_c} - p_2}{2C} \right) V_1(p_1^*, p_2^*). \end{aligned} \quad (23)$$

The first term in equation (23) reflects the fact that a consumer whose type is less than the marginal type will buy from firm 1 in the current period and will be locked in with firm 1 at the same price in all subsequent periods. The second term reflects the fact that consumers with types above the marginal type do not buy from firm 1 today and will still be free buyers tomorrow. Similar considerations are at work in the definition of firm 2's value function:

$$V_2(p_1, p_2) = (p_2 + \beta_f V_2(p_1^*, p_2^*)) \left(\frac{1}{2} + \frac{u_2^* - u_1^*}{2C} + \frac{\frac{p_1}{1-\beta_c} - p_2}{2C} \right). \quad (24)$$

Put $v_i^* = V_i(p_1^*, p_2^*)$, $i = 1, 2$. Then, the equilibrium value functions solve the following equations:

$$v_1^* = \frac{1}{1-\beta_f} p_1^* x^* + \beta_f (1-x^*) v_1^*, \quad (25)$$

$$v_2^* = (p_2^* + \beta_f v_2^*) (1-x^*). \quad (26)$$

Taking first-order conditions in equations (23) and (24) and plugging in $p_1 = p_1^*$ and $p_2 = p_2^*$, we get:

$$0 = \frac{1}{1 - \beta_f} \left(x^* - \frac{1}{2C} \frac{p_1^*}{1 - \beta_c} \right) + \frac{1}{2C(1 - \beta_c)} \beta_f v_1^*, \quad (27)$$

$$0 = 1 - x^* - \frac{1}{2C} (p_2^* + \beta_f v_2^*). \quad (28)$$

To summarize, if (p_1^*, p_2^*) is a stationary subgame-perfect equilibrium, then $(u_1^*, u_2^*, x^*, v_1^*, v_2^*, p_1^*, p_2^*)$ jointly solves equations (19), (21), (22), (25), (26), (27), and (28), and $x^* \in [0, 1]$. We show that this system of equations has a unique solution, and that this solution is indeed a stationary subgame-perfect equilibrium.

Lemma 3. *There exists a unique stationary subgame-perfect equilibrium in the mixed subgame.*

Proof. From now on, we drop the star superscripts to ease notation. The proof is analytical, but some of the computations are cumbersome. Details of the calculations can be found in Mathematica file `ETF-Hotelling.nb`.

We approach the problem as follows. Fix some $x \in [0, 1]$. Using Mathematica, we show that there exists a unique vector $(u_1(x), u_2(x), v_1(x), v_2(x), p_1(x), p_2(x))$, which jointly solves equations (19), (22), (25), (26), (27), and (28) for this value of x (Step 1 in the Mathematica file). In Step 2, we plug this vector into equilibrium condition (21) and show that the condition holds if and only if x is a root of polynomial

$$P(X) = -(1 - \beta_f)(3 - \beta_c - 2\beta_f) + 2(1 - \beta_f)(3 - 2\beta_c - 2\beta_f)X + (\beta_c(1 - 3\beta_f) + 2\beta_f(2 - \beta_f))X^2.$$

If $\beta_c = \beta_f = 0$, then P is linear. Therefore, P has a unique root: $\hat{x} = 1/2$. If $\beta_c \neq 0$ or $\beta_f \neq 0$, then P is quadratic. We show that P is convex, $P(0) < 0$ and $P(1) > 0$ (Step 3). Therefore, there exists a unique $\hat{x} \in [0, 1]$ such that $P(\hat{x}) = 0$. This \hat{x} is the highest root of quadratic polynomial P . Therefore, there exists at most one stationary subgame-perfect equilibrium. Conversely, since the objective functions in (23) and (24) are strictly concave in p_1 and p_2 , respectively, first-order conditions are sufficient for optimality. Therefore, the profile of prices pinned down by \hat{x} is a stationary subgame-perfect equilibrium, and the mixed subgame has a unique stationary subgame-perfect equilibrium.

\hat{x} , $v_1(\hat{x})$, and $v_2(\hat{x})$ are computed in Step 4. Recall that v_i is the value of firm i per free buyer. At stage 0, there is a mass 1 of free buyers. Therefore, v_i also gives us the present discounted value of firm i 's profit flows at time 0. \square

Appendix A.3 The ETF game

Now that we have fully characterized the equilibria in all subgames starting at stage 0, we can go back to stage -1 and solve for the equilibria of the ETF game. In stage -1, firms choose between actions n (no ETF) and e (ETF). In the following, we let π_k^l be the equilibrium profit of a firm that plays action $k \in \{n, e\}$ when its rival plays action $l \in \{n, e\}$. By Lemmas 1, 2, and 3, $v_n^n = \frac{1}{1 - \beta_f} \frac{C}{2}$, $v_e^e = \frac{1 - \beta_c}{1 - \beta_f} \frac{C}{2}$, and v_e^n and v_n^e are solved for in Appendix A.2.3.

All we need to do now is solve the following two-by-two game:

	n	e
n	(v_n^n, v_n^n)	(v_n^e, v_n^n)
e	(v_e^n, v_e^n)	(v_n^n, v_n^n)

We prove the following lemma:

Lemma 4. *There exist functions $\beta_c^n, \beta_c^e : [0, 1) \rightarrow [0, 1)$ such that for every $(\beta_c, \beta_f) \in [0, 1)^2$,*

- $v_n^n \geq v_e^n$ (resp. $v_n^n \leq v_e^n$) if and only if $\beta_c \geq \beta_c^n(\beta_f)$ (resp. $\beta_c \leq \beta_c^n(\beta_f)$),
- $v_e^e \geq v_n^e$ (resp. $v_e^e \leq v_n^e$) if and only if $\beta_c \leq \beta_c^e(\beta_f)$ (resp. $\beta_c \geq \beta_c^e(\beta_f)$).

In addition, $\beta_c^n(0) = \beta_c^e(0) = 0$, and $\beta_c^n(\beta_f) < \beta_c^e(\beta_f)$ for every $\beta_f > 0$.

Proof. We normalize C to 1 without loss of generality. Let us first study expression $v_n^n - v_e^n$. It is straightforward to check that, when $\beta_f = \beta_c = 0$, $v_e^n = \frac{1}{2} = v_n^n$, and that $v_n^n > v_e^n$ whenever $\beta_c > \beta_f = 0$ (Step 5 in the Mathematica file). Therefore, $\beta_c^n(0) = 0$. In the following, we assume that $\beta_f > 0$. For every $\beta_f \in (0, 1)$, define the following quartic polynomial:

$$P_{\beta_f}^n(X) = (48\beta_f - 64\beta_f^2 + 16\beta_f^3 + 4\beta_f^4) + (-72 + 16\beta_f + 52\beta_f^2 - 20\beta_f^3) X \\ + (129 - 86\beta_f + 9\beta_f^2) X^2 + (-72 + 24\beta_f) X^3 + 16X^4.$$

Using Mathematica (Step 6), we show that $v_n^n \geq v_e^n$ (resp. $v_n^n \leq v_e^n$) if and only if $\beta_c \geq \beta_c^n(\beta_f)$ (resp. $\beta_c \leq \beta_c^n(\beta_f)$), where $\beta_c^n(\beta_f)$ is equal to the first real root of polynomial $P_{\beta_f}^n(X)$ if $\beta_f \leq 9/10$, and to the second real root of the same polynomial if $\beta_f > 9/10$.

Next, we turn our attention to the expression $v_e^e - v_n^e$. As before, it is straightforward to check that when $\beta_f = \beta_c = 0$, $v_n^e = \frac{1}{2} = v_e^e$, and that $v_e^e < v_n^e$ whenever $\beta_c > \beta_f = 0$ (see Step 5). Therefore, $\beta_c^e(0) = 0$. In the following, we assume that $\beta_f > 0$. For every $\beta_f \in (0, 1)$, define the following polynomials:

$$P_{\beta_f}^e(X) = (-48\beta_f + 80\beta_f^2 - 32\beta_f^3 + 4\beta_f^4) + (72 - 80\beta_f + 12\beta_f^2 - 28\beta_f^3 + 8\beta_f^4) X \\ + (-87 + 98\beta_f + 9\beta_f^2 + 4\beta_f^4) X^2 + (30 - 28\beta_f - 14\beta_f^2 - 4\beta_f^3) X^3 + (1 - 6\beta_f + 9\beta_f^2) X^4, \\ Q(X) = 167 + 926X + 1079X^2 + 384X^3.$$

Define $\beta_c^e(\beta_f)$ as the second real root of polynomial $P_{\beta_f}^e(X)$ when $\beta_f \neq 1/3$, and as minus the first real root of polynomial $Q(X)$ when $\beta_f = 1/3$. In Step 7, we show that $v_e^e \geq v_n^e$ (resp. $v_e^e \leq v_n^e$) if and only if $\beta_c \leq \beta_c^e(\beta_f)$ (resp. $\beta_c \geq \beta_c^e(\beta_f)$).

Finally, we show that inequalities $v_n^n \leq v_e^n$ and $v_e^e \leq v_n^e$ cannot hold simultaneously when $\beta_f > 0$ (Step 8). It follows that $\beta_c^n(\beta_f) < \beta_c^e(\beta_f)$ for every $\beta_f > 0$. \square

We can proceed with the complete characterization of equilibria.

Proposition 1. *Let $(\beta_c, \beta_f) \in [0, 1)^2$. Then:*

- (n, n) is an equilibrium if and only if $\beta_c \geq \beta_c^n(\beta_f)$.
- (e, e) is an equilibrium if and only if $\beta_c \leq \beta_c^e(\beta_f)$.
- (n, e) and (e, n) are equilibria if and only if $\beta_f = \beta_c = 0$.

For every $\beta_f > 0$, the set of β_c s such that equilibria (n, n) and (e, e) coexist is an interval with a non-empty interior.

Proof. This follows immediately from Lemma 4. □

Figure 10: Equilibrium characterization




Figure 10 provides a graphical representation of Proposition 1 in the (β_f, β_c) plane. The white curve is threshold $\beta_c^n(\cdot)$. The black curve is threshold $\beta_c^e(\cdot)$. (n, n) (resp. (e, e)) is an equilibrium if and only if (β_f, β_c) lies above the white curve (resp. below the black curve). In the dark shaded area, equilibria (n, n) and (e, e) coexist, and firms face a coordination problem. They might end up playing the (e, e) equilibrium, which, as discussed in Appendix A.2.2, yields lower payoffs and worse allocative efficiency than the (n, n) equilibrium.

Appendix B Additional results for different sets of instrumental variables

Table 16: Handset-carrier fixed-effect regressions, 16,408 observations

parameters	unrestricted		restricted				
	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV	(7) IV*
handset price, P_{jt} (s.e.)	-4.568 (0.263)	-21.033 (16.482)	-4.591 (0.262)	-15.565 (14.890)			
service fee, p_{jt} (s.e.)	1.085 (0.091)	-19.084 (5.420)	1.021 (0.091)	-17.273 (4.345)			
total cost, $p_{jt} + P_{jt}$ (s.e.)					0.340 (0.085)	-17.190 (4.296)	-19.185 (6.030)
constant (s.e.)	-6.632 (0.066)	2.945 (2.719)	-8.070 (0.046)	1.702 (2.216)	-8.148 (0.046)	1.787 (2.092)	2.757 (2.935)
product fixed effect	yes	yes	yes	yes	yes	yes	yes
time fixed effect	yes	yes	no	no	no	no	no
carrier-time fixed effect	no	no	yes	yes	yes	yes	yes
<i>first stage statistics</i>							
F statistic, P_{jt} (p-value)		44.68 (0.000)		11.76 (0.000)			
F statistic, p_{jt} (p-value)		54.20 (0.000)		14.34 (0.000)			
F statistic, $P_{jt} + p_{jt}$ (p-value)						18.61 (0.000)	18.99 (0.000)

Notes: Regressions labeled “IV” are two-stage least squares with 4 instrumental variables; regression “IV*” has only 2 instruments: average age and consumer satisfaction by products of competitors.

Table 17: Handset fixed-effect regressions, 16,408 observations

parameters	unrestricted		restricted		
	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) IV*
handset price, P_{jt} (s.e.)	-3.549 (0.284)	-4.592 (9.098)			
service fee, p_{jt} (s.e.)	1.397 (0.102)	-10.030 (2.968)			
total cost, $p_{jt} + P_{jt}$ (s.e.)			0.755 (0.094)	-9.574 (2.846)	-9.121 (3.195)
constant (s.e.)		1.191 (1.437)	-3.640 (0.173)	1.368 (1.396)	1.149 (1.563)
handset fixed effect	yes	yes	yes	yes	yes
carrier-time fixed effect	yes	yes	yes	yes	yes
R ²	0.65	0.35	0.65	0.36	0.38
<i>first stage statistics</i>					
F statistic, P_{jt} (p-value)		5.98 (0.000)			
F statistic, p_{jt} (p-value)		7.22 (0.000)			
F statistic, $P_{jt} + p_{jt}$ (p-value)				6.58 (0.000)	10.03 (0.000)

Notes: Regressions labeled “IV” are two-stage least squares with 4 instrumental variables; regression “IV*” has only 2 instruments: average age and consumer satisfaction by products of competitors.

The overidentifying restrictions test does not reject any of the specifications in Table 3. All parameter estimates are significant at the 1 percent level. Due to the relatively large standard errors, the t-test rejects a statistically significant difference in the parameter

estimates across specifications with alternative sets of IVs. In particular, when comparing the difference between (1) and (2) the t-statistic has a p-value of 0.69; when comparing (1) and (3) the p-value is 0.42; the p-value of the test for difference in coefficients between (2) and (3) is 0.65.¹⁶

To select our main specification we use a different criterion, namely the share of products with negative predictions for marginal costs. It turns out that parameter estimates in specification (3) result in the smallest share of negative marginal cost predictions among all specifications (see Table 10 for further details). We used this criterion to choose our preferred specification because estimation of the dynamic demand model in our case does not rely on any supply-side moment conditions that, if included, wouldn't admit negative cost estimates. In what follows, in the main text we present an extended analysis of the results based on the estimates from specification (3). The results from other specifications are mentioned briefly whenever relevant with the figures and tables reported in appendices.

Table 18: Second-stage optimal GMM parameter estimates and elasticity predictions

parameter	estimates		
	(1)	(2)	(3)
price coefficient, α_p	-5.428	-6.466	-8.163
(s.e.)	(1.496)	(2.155)	(3.014)
carrier-time fixed effects	yes	yes	yes
handset fixed effects	yes	yes	yes
<i>service fee elasticity</i>			
average	-1.973	-2.350	-2.967
median	-2.017	-2.403	-3.033
standard deviation	0.557	0.663	0.838
<i>handset price elasticity</i>			
average	-0.350	-0.416	-0.524
median	-0.274	-0.326	-0.410
standard deviation	0.270	0.321	0.404
Hansen's J-stat	3.033	2.164	0.841
(p-value)	(0.386)	(0.339)	(0.359)

Notes: (1) employs IV1 through IV4; (2) uses IV1, IV2, IV3; instruments for specification (3) include IV2 and IV3 only. Estimation results from other combinations of instrumental variables are available upon request.

¹⁶The t-test was computed assuming zero covariance between the coefficients.

Figure 11: Distribution of own price elasticity for all products, specification (2)




Figure 12: Distribution of own price elasticity for all products, specification (1)




Table 19: Changes in consumer welfare and market shares, one-type model (2)

counterfactual scenario		mean	p50	min	max	sd
ETFs	handset	change in value functions				
No	purchased at obs. prices	0.56	0.54	0.49	0.70	0.04
No	purchased at new prices	0.34	0.33	0.26	0.48	0.03
No	rented	0.84	0.82	0.74	1.01	0.04
Yes	rented	0.14	0.14	0.12	0.16	0.01
ETFs	handset	change in market shares				
No	purchased at obs. prices	0.40	0.37	-0.67	1.60	0.27
No	purchased at new prices	0.55	0.56	-0.81	2.75	0.46
No	rented	0.57	0.39	-0.79	10.40	0.74
Yes	rented	0.23	0.10	-0.51	4.77	0.47

Notes: Service fees are fixed at observed levels; changes are calculated using $\Delta(v)_i = (v(i) - V(0))/V(0)$, where i denotes scenario, $v(i)$ is variable of interest and $V(0)$ is factual outcome.

Table 20: Changes in consumer welfare and market shares, one-type model (1)

counterfactual scenario		mean	p50	min	max	sd
ETFs	handset	change in value functions				
No	purchased at obs. prices	0.44	0.43	0.39	0.55	0.03
No	purchased at new prices	0.27	0.26	0.20	0.38	0.03
No	rented	0.66	0.64	0.59	0.78	0.03
Yes	rented	0.12	0.12	0.10	0.13	0.01
ETFs	handset	change in market shares				
No	purchased at obs. prices	0.35	0.33	-0.59	1.27	0.21
No	purchased at new prices	0.49	0.50	-0.74	2.14	0.37
No	rented	0.50	0.36	-0.72	7.11	0.57
Yes	rented	0.19	0.09	-0.45	3.38	0.36

Notes: Service fees are fixed at observed levels; changes are calculated using $\Delta(v)_i = (v(i) - V(0))/V(0)$, where i denotes scenario, $v(i)$ is variable of interest and $V(0)$ is factual outcome.

Table 21: Changes in consumer welfare and market shares, four-type model

counterfactual scenario		mean	p50	min	max	sd
ETFs	handset	change in value functions				
No	purchased at obs. prices	0.90	1.05	0.22	1.86	0.39
No	purchased at new prices	0.58	0.69	0.15	1.29	0.26
No	rented	1.41	1.69	0.30	2.98	0.64
Yes	rented	0.22	0.26	0.08	0.35	0.09
ETFs	handset	change in market shares				
No	purchased at obs. prices	0.81	0.51	-0.84	10.74	1.09
No	purchased at new prices	1.01	0.51	-0.95	20.64	1.60
No	rented	1.48	0.46	-0.93	235.87	3.89
Yes	rented	0.56	0.15	-0.71	41.96	1.56

Notes: Service fees are fixed at observed levels; changes are calculated using $\Delta(v)_i = (v(i) - V(0))/V(0)$, where i denotes scenario, $v(i)$ is variable of interest and $V(0)$ is factual outcome.

Table 22: Change in service fees offsetting consumer gains from ETF elimination, %

type of compensating change	spec.(3)	spec.(2)	spec.(1)
	$\alpha_p = -8.16$	$\alpha_p = -6.47$	$\alpha_p = -5.43$
increase in service fees at obs. h-set prices	42.59	43.09	43.38
increase in service fees at new h-set prices	31.70	31.36	31.12

Notes: Offsetting price increases are computed such that the differences between consumer value functions before the ETF elimination and consumer value functions after the ETF elimination and corresponding change in service fees are zero on average.