

Estimating Dynamic Demand in Differentiated Durable Goods Markets Using Price Alert Thresholds

Serafin Grundl*

January 11, 2014

Abstract

Retail prices are volatile and decrease over time in many markets for differentiated durable goods. Consumers in these markets have an incentive to delay their purchase and wait for lower prices. This paper uses a novel data set from a price alert service for TVs sold on Amazon.com to estimate substitution patterns across products and over time. Users of a price alert service submit a price threshold for a product they want to purchase and receive an alert when their threshold is reached. I estimate consumer preferences using a discrete/continuous-choice model in which the price threshold is the solution to an optimal stopping problem. The estimates imply that many consumers respond to a price increase by delaying their purchase rather than by choosing a different TV. This suggests that TVs mostly compete with their own future selves rather than with other TVs. To measure how much consumers benefit from delaying their purchase I consider a counterfactual in which consumers buy a TV immediately upon arrival at the market.

*I am deeply indebted to my advisors Amit Gandhi, Ken Hendricks and Alan Sorensen for their guidance and support. I also thank Rasmus Lentz, Jack Porter, Dan Quint and Yu Zhu for their helpful comments. I am very grateful to Dan Green and the team of camelcamelcamel.com for providing the data which made this project possible. All remaining errors are my own. Contact: Department of Economics, University of Wisconsin-Madison, 7310 Social Science, 1180 Observatory Dr., Madison, WI 53706, grundl@wisc.edu.

1 Introduction

In many markets for differentiated durable goods consumers face decreasing and volatile prices. Price decreases are often driven by decreasing marginal costs. Well known examples include technology products with CPUs or LCDs, for which learning by doing leads to decreasing marginal costs. Even if marginal costs remain constant, prices can decrease over time if firms engage in intertemporal price discrimination and try to sell to consumers with high valuations first and later at lower prices to consumers with lower valuations. While retail prices tend to decrease in the long run, price increases are common and retail prices are volatile, especially online. Online retailers such as Amazon.com use algorithmic pricing tools to change their prices frequently. Depending on the product category, prices can change several times a week or even several times a day.¹

Consumers in these markets must decide which product to purchase and when to purchase. The downward trend of prices and their volatility create an option value of waiting, and give consumers an incentive to delay their purchase and wait for better prices. Consumers can respond to a price increase of a product either by purchasing a different product or by purchasing the same product at a later time. Estimating these substitution patterns across products and over time is crucial to understand whether a product competes mostly with other products or with its own future selves.

This paper studies a data set from a price alert service, which sheds light on how consumers plan their purchase in a market for differentiated durable goods with volatile prices. A price alert service helps consumers to track the prices of online retailers. The consumer can submit a price threshold for a product and receives an alert when the threshold is reached. The data comes from the price alert service of camelcamelcamel.com. It contains price thresholds for TVs sold on Amazon which were submitted by approximately 10,000 users between May 2008 and May 2012 and price histories during this time period. A limitation of the data set is that the purchase decision is not observed.

As the majority of consumers submits a price threshold for a single TV, I propose a tractable model in which the consumer keeps track of a single price.² The consumer in the model solves an optimal stopping problem in which she waits for the right price to purchase the TV she tracks. The solution of the optimal stopping problem takes the form of a threshold rule and can therefore be implemented with the help of a price alert service. When the consumer arrives at the market, she selects the TV for which the optimal stopping problem is associated with the highest expected utility. The model is tractable because the

¹This practice has received attention in several newspapers. For example: The New York Time from 11/30/12: “Retail Frenzy: Prices on the Web Change Hourly”, The Wall Street Journal from 9/5/12: “Coming Soon: Toilet Paper Priced Like Airline Tickets”, Financial Times from 7/8/12: “Amazon ‘robo-pricing’ sparks fears”.

²The discussion in section 3 lays out a modelling approach for consumers who track prices of multiple TVs and explains why estimation of the model for these consumers is computationally burdensome.

decisions of what to purchase and when to purchase are made in two steps and therefore only the price of the selected TV enters the optimal stopping problem as a state variable.

The core of the paper is devoted to estimation of consumer preferences. Consumers are characterised by a vector of valuations for different TVs. Under the assumption that the discount rate is known, the joint distribution of valuations is shown to be nonparametrically identified through price variation faced by consumers who arrive at the market at different times. For estimation purposes, preferences are specified as a function of product characteristics which are associated with random coefficients. Estimation proceeds in two steps. In the first step, the optimal price threshold function and the associated value function are estimated. The optimal threshold function maps a valuation for a TV into the optimal price alert threshold. The associated value function maps a valuation for a TV and its current price into the expected utility obtained from submitting the optimal threshold. Both functions depend on the stochastic process which governs the evolution of prices and are estimated using data on price histories. In the second step consumer choices of TVs and price thresholds are used to estimate preferences for TV characteristics.

Substitution across TVs and over time are quantified by simulating a price increase for a TV. To summarize the substitution patterns I consider the consumers who would have bought the TV within one month of their arrival at the market, but no longer do so in response to the price increase. These consumers respond to the price increase either by choosing a different TV, or by delaying their purchase of the same TV to a later time. The fraction of consumers who substitutes to a later time, rather than a different TV, ranges from 71% for inexpensive entry models to 89% for expensive high end models. This suggests that TVs mostly compete with their own future selves rather than with other TVs.

To understand how much consumers gain from delaying their purchase I consider a counterfactual Buy Now scenario in which consumers buy a TV immediately upon arrival at the market. Tracking the price of a single TV to delay the purchase increases average consumer surplus by \$104 or 1.7% compared to the Buy Now counterfactual. To put the gain of \$104 into perspective I consider the benchmark of perfect foresight for all TVs which serves as an upper bound on consumer surplus.³ I find that the gain of \$104 corresponds to 60% of the gains under perfect foresight for all TVs. Therefore consumers realize a fairly large fraction of the potential gains from delaying their purchase by tracking the price of a single TV. This could explain why most consumers choose to submit a threshold for a single TV.

The Buy Now scenario is also useful to understand how intertemporal substitution affects competition among TVs. Own-price elasticities and market shares differ considerably from the model with intertemporal substitution. As the Buy Now scenario does not consider

³Ideally we would like to consider the surplus of a consumer who tracks the prices of all TVs but does not have perfect foresight. Unfortunately this decision problem is not tractable because the state space includes the prices of all TVs.

the possibility of intertemporal substitution it predicts more substitution across TVs. The predicted own-price elasticities for all TVs are larger than in the model with intertemporal substitution - for some high-end TVs by more than 100%. Market shares in the Buy Now scenario are higher for entry models (by up to 44%) and lower for high end models (by up to 49%). The reason for this shift in market shares is that delaying the purchase is more valuable for expensive high-end TVs, because they exhibit larger absolute price jumps.

Related Literature

There is a large literature estimating static demand models for differentiated goods (e.g. Berry, Levinsohn and Pakes (1995)). Static demand estimates have been used for example to study the effect of mergers (Nevo (2000)) or to examine the benefits from new products (Petrin (2002)). Many applications in this literature have studied markets for durable goods.

Intertemporal substitution can limit the market power of firms in markets for durable goods. Coase (1972) conjectured that the market power of a monopolist selling a homogeneous good would be eliminated through competition with its own future selves and sparked a large theoretical literature on this issue (e.g. Stokey (1979) and Bulow (1982)).

This paper adds to the literature estimating dynamic demand for differentiated durable goods. Most of the existing literature uses purchase data which is aggregated over retailers, consumers and time, typically a panel of prices and market shares which is measured monthly or quarterly. The model framework combines features from static demand models for differentiated goods (Berry, Levinsohn and Pakes (1995)) with the dynamic discrete choice framework for homogenous goods of Rust (1987). Melnikov (2013) first estimated such a model by introducing a simplifying assumption to reduce the dimensionality of the state space.⁴ Several authors have estimated similar models, notably Song and Chintagunta (2003), Nair (2007) and Carranza (2007). The state of the art in the literature is well exemplified by Gowrisankaran and Rysman (2012) who add random coefficients and replacement purchases to Melnikov's framework. Conlon (2010) demonstrates how to simplify the estimation of a variant of Gowrisankaran and Rysman's model using a MPEC method.⁵ Lee (forthcoming) and Schiraldi (2011) introduced complementarities and a second hand market, respectively. Outside of the durable good context Shcherbakov (2013) and Nosal (2012) use the framework to estimate models with switching costs. Similar models are also used to estimate demand for storable goods in Hendel and Nevo (2006) and Erdem, Imai and Keane (2003) using consumer level data.

One important difference between this paper and the existing durable goods literature is that I consider intertemporal decisions over relatively short horizons (weeks or months) for which the volatility of prices plays an important role. Aggregation over time masks the

⁴An earlier version of this paper was circulated in 2000.

⁵MPEC stand for Mathematical Programming with Equilibrium Constraints, see also Su and Judd (2012).

volatility of prices and is not well suited to study intertemporal decisions over short horizons. The literature studies intertemporal decisions over longer horizons and therefore considers additional aspects such as replacement purchases, the anticipation of new products arriving at the market and time-varying unobservable product characteristics.⁶

The literature assumes that consumers check the prices of all goods in every time period and decide whether to buy one of the available goods or to wait. This is a natural assumption if purchase data is observed. Without making further assumptions, however, this decision problem is typically untractable because the state space includes the prices of all products.⁷ Two simplifying assumptions to reduce the dimensionality of the state space have been explored. The first assumption is perfect foresight, i.e. consumers know the future prices. Under perfect foresight time is the only state variable. The second assumption is called inclusive value sufficiency. The logit inclusive value arises from the type I extreme value error in the dynamic discrete choice framework. The state variables enter the decision problem only through their impact on the logit inclusive value. The dimensionality of the state space is reduced by assuming that the current value of the inclusive value is sufficient to predict the inclusive value in the next period. See Gowrisankaran and Rysman (2012) for a more detailed discussion of the inclusive value sufficiency assumption.

In this paper the dimensionality of the state space is reduced in a different fashion which is motivated by the data. As the majority of consumers submits a threshold for a single TV I consider a simpler decision problem in which the decisions of what to purchase and when to purchase are made in two steps. Therefore only the price of the selected TV enters the optimal stopping problem as a state variable.

Preferences are assumed to be stable over time in this paper. The dynamic discrete choice framework includes preference shocks which are iid across consumers, products and also over time. This has an effect on the substitution patterns predicted by the model. The consumer is more likely to switch products if the purchase is delayed due to a new draw of preference shocks. For example Conlon (2010), who also studies demand for TVs, concludes in contrast to this paper that substitution towards delaying the purchase of the same product is fairly unimportant.

The remainder of the paper is organized as follows. Section 2 contains a description of the data. The model is introduced in section 3. Section 4 discusses identification. A two step estimation procedure is proposed in section 5. Section 6 presents the results and discusses robustness issues. Section 7 concludes.

⁶Conlon (2010) argues that replacement purchases for large TVs are rare. As I estimate the model with six months of data it is unlikely that consumers replace their TV during the sample period. New TVs are usually introduced in spring. To focus on price changes for existing models rather than the introduction of new models the analysis in this paper is restricted to six months from August to January when only few new TVs are introduced.

⁷If there are other time-varying product characteristics or the set of products is changing over time the state space is even larger.

2 Data Description

The data set contains prices for TVs sold on Amazon.com and price thresholds submitted on camelcamelcamel.com between May 2008 and May 2012. Over the whole sample period there are 10013 users who submitted a price threshold for a TV. A limitation of the data set is that purchase decisions are not observed.

TV prices on Amazon.com are volatile and tend to decrease. An example of a price path is shown in Figure 1.



Figure 1: Price path for the 46 inch Samsung UN46D7000 from March 2011 to March 2012. Screenshot from camelcamelcamel.com.

Prices change on average twice a week and the average price change is -0.2% . Figure 2 presents the distribution of price changes. While most price changes are small, 13.6% of the price changes are larger than 10% in magnitude.

Figure 3 shows undiscounted savings under perfect foresight over a 1 month decision horizon as a function of the initial TV price. The decision horizon restricts how many weeks after the initial price was posted are taken into consideration. The figure shows that due to the volatility of prices delaying the purchase can result in substantial savings even over short horizons. This suggests that intertemporal substitution over short horizons is important in this environment. Potential savings are larger for expensive TVs which exhibit larger absolute price changes.

A striking feature of the data is that most users track only few TVs. 64% of the users submit a price threshold for a single TV. 18% and 7% of the users submit thresholds for two or three TVs respectively. The remaining 11% follow four or more TVs.

Users can submit more than one price threshold for the same TV if it is sold by multiple




Figure 2: Distribution of price changes for new TVs. Changes which are larger than 10% in magnitude are collected are in the two outside bars. The average price change is -0.2% and the median price change is -0.14%.

retailers on Amazon.com. Camelcamelcamel.com keeps track of three different price histories: Amazon’s own price for a new TV, the price of a 3rd party retailer for a new TV and the price of a 3rd party retailer for a used TV.⁸ 81% of the submitted thresholds track Amazon’s own price, 17% and 2% track the 3rd party price for a new and used TV, respectively. As the prices charged by Amazon and the other retailers for new TVs are highly correlated tracking more than one price rarely results in significant additional savings.

The average price threshold is set \$173 or 10.8% below the current price. The distribution of the gap between the price and the price threshold is shown in Figure 4 (in percentage terms).⁹ The median waiting time before the threshold is reached is 9 weeks. Prices change discontinuously which results in overshoot, i.e. the price when the threshold is reached can be significantly lower than the threshold. The average overshoot when the threshold is reached is \$52 or 3.7% of the price threshold.

Some users choose their price threshold from a list of predefined threshold prices. The list contains discounts of 1 cent, 10%, 25%, 50%, 75% and 90%. 18% of the submitted thresholds are 1 cent below the current price while the other predefined options play a less important role.

The users seldom update their price thresholds. 87% of the thresholds are never changed.

⁸Users cannot see the 3rd party retailer whose price is displayed. Usually the lowest 3rd party prices are displayed.

⁹The distribution of the gap between the price and the price threshold in absolute terms can be found in Figure 8 in Appendix D.




Figure 3: Undiscounted savings under perfect foresight over a 1 month decision horizon as a function of the initial TV price. The decision horizon restricts how many weeks after the initial price was posted are taken into consideration. The unit of observation is a TV/day pair. The graph shows a local linear regression of savings on the price using TVs between 32 and 65 inches screen size which are sold directly by Amazon.

The thresholds which are changed can be divided into two groups: 4% are changed before the first price change and are often corrections of obvious mistakes such as a missing digit. 9% are changed later and can be considered 'real revisions'.¹⁰

¹⁰These revisions occur with approximately equal probability before (49%) and after (51%) the threshold has been reached. 59% of the revisions which occur before the threshold has been reached increase the threshold. As expected almost all revisions which occur after the threshold has been reached (92%) decrease the threshold. The remaining 8% are typically consumers who missed the time window when the price was lower than their threshold and later return to submit a less aggressive threshold.




Figure 4: Distribution of the gap between the price and the submitted threshold as a fraction of the price.

3 Model

Consider consumer i who arrives at the market at time t . The consumer is characterized by nonnegative valuations v_{ij} for the TVs $j = 1 \dots J$. The valuations are the lifetime utility the consumer gets from owning the TV. Valuations are drawn from a distribution with density f on the support $[v_1, \bar{v}_1] \times \dots \times [v_J, \bar{v}_J]$ with $v_1, \dots, v_J \geq 0$. Time is continuous and discounted at a rate ρ , which is shared by all consumers.

The model applies to consumers who submit a single price threshold. When the consumer arrives at the market she faces a discrete/continuous choice problem in which she has to choose a TV and a price threshold. The threshold is the solution to an optimal stopping problem in which the consumer waits for the right price to purchase the TV. The consumer selects the TV for which the optimal stopping problem is associated with the highest expected utility.

Threshold Choice

Consider a consumer who constantly monitors the price of TV j . At any time $t + T$ the consumer can stop and purchase the TV which yields

$$\exp(-\rho T) (v_{ij} - P_{jt+T}).$$

To solve this optimal stopping problem we must specify a process for the price of TV j .

Assumption 1 (Price Process). *The price of TV j evolves according to*

$$P_{jt} = P_{j0} \exp(X_{jt}),$$

where X_{jt} follows a compound Poisson Process with arrival rate λ and jump size distribution G which is started at $X_{j0} = 0$.

This jump process can capture that prices change discontinuously and decrease in a stochastic fashion. The process also reflects that potential savings are larger for expensive TVs because all TVs share the same distribution of relative price changes G , and expensive TVs experience therefore larger absolute price changes. A limitation is that this specification does not capture 'sales' where prices tend to return to their previous level after they have been temporarily reduced. The specification makes the decision problem tractable because neither the past and current prices of TV j nor those of other TVs help to predict changes of $\log(P_{jt})$.

The optimal stopping problem of the consumer is identical to the problem of an investor who waits for the optimal time to exercise a perpetual American put option in which v_{ij} takes the role of the strike price. Mordecki (2002) showed that this problem is solved by a threshold rule if the price follows an exponential Levy process such as the process defined in Assumption 1.¹¹ The optimal stopping policy can therefore be implemented with the help of a price alert service. Mordecki shows that the optimal threshold has the following form:

$$\bar{P}(v_{ij}, \rho) = v_{ij} \mathbb{E}[Z] = v_{ij} s(\rho), \quad (1)$$

where

$$Z = \exp\left(\inf_{0 \leq t' \leq \tau(\rho)} X_{jt'}\right),$$

where $\tau(\rho)$ is an exponential random variable with mean $\frac{1}{\rho}$. In words, Z is generated by stopping the infimum process of the normalized price $\frac{P_{jt}}{P_{j0}}$ at a random time $\tau(\rho)$. If ρ increases the process tends to be stopped earlier and therefore s is increasing in ρ . As Z takes on values between zero and one the optimal threshold is a fraction of the consumer's valuation for the TV.

¹¹Mordecki was the first to show that the result holds for all exponential Levy processes. Several other authors have shown the result under some restrictions. For references see Alili and Kyprianou (2005) and Mordecki (2002).

TV Choice

Let $W(v_{ij}, P_{jt}, \rho)$ be the expected utility from submitting the optimal threshold for TV j . The user submits a threshold for the TV which offers the highest expected utility

$$\arg \max_j W(v_{ij}, P_{jt}, \rho).$$

There is no outside option required because $W(v_{ij}, P_{jt}, \rho)$ is nonnegative for all valuations and prices. Mordecki shows that W takes the following form:

$$W(v_{ij}, P_{jt}, \rho) = (v_{ij} - \bar{P}(v_{ij}, \rho)) \widetilde{W}\left(\frac{P_{jt}}{\bar{P}(v_{ij}, \rho)}, \rho\right),$$

where

$$\widetilde{W}\left(\frac{P_{jt}}{\bar{P}(v_{ij}, \rho)}, \rho\right) = \mathbb{E} \left[\max \left\{ \left(1 - \frac{P_{jt}Z}{\bar{P}(v_{ij}, \rho)}\right) / (1 - s(\rho)), 0 \right\} \right]. \quad (2)$$

It is easy to see that W is increasing in v_{ij} and decreasing in P_{jt} and ρ .

Remark 1. There is no closed form expression for \widetilde{W} because it depends on the joint distribution of the purchase time and the purchase price for a given price threshold. However, \widetilde{W} can be estimated from price data as will be explained in section 5. It turns out that the estimated \widetilde{W} can be closely approximated as follows:

$$\widetilde{W}\left(\frac{P_{jt}}{\bar{P}(v_{ij}, \rho)}, \rho\right) \approx \left(\frac{P_{jt}}{\bar{P}(v_{ij}, \rho)}\right)^{g(\rho)}, \quad (3)$$

for some $g(\rho) < 0$ and $g(\rho)$ increasing.¹² The approximation follows the estimated \widetilde{W} very closely except for very large values of the arguments which are not relevant in practice.¹³ For details about the approximation of \widetilde{W} refer to Appendix A.

This approximation provides some further insight into the trade-off which determines the TV choice. Using the approximation the TV choice can be reduced to a linear discrete choice model:

¹²This functional form of \widetilde{W} is a generalization of the case where the price of TV j follows a geometric Brownian motion where g depends on the percentage drift and volatility.

¹³The maximum approximation error for an annual discount factor of 0.9 and $4\bar{P}(v_{ij}, \rho) > P_{jt}$, for example, is less than 0.02%.

$$\begin{aligned}
& \arg \max_j W(v_{ij}, P_{jt}, \rho) \\
&= \arg \max_j v_{ij} (1 - s(\rho)) \left(\frac{P_{jt}}{v_{ij} s(\rho)} \right)^{g(\rho)} \\
&= \arg \max_j \log(v_{ij}) + \frac{g(\rho)}{1 - g(\rho)} \log(P_{jt})
\end{aligned}$$

The consumer trades off the valuation for the TV $\log(v_{ij})$ and the disutility from waiting $\frac{g(\rho)}{1-g(\rho)} \log(P_{jt})$. To understand why the second term represents disutility from waiting, remember that the current price of TV j does not affect the optimal price threshold, but it does affect how long the consumer has to wait until the threshold is reached. As ρ increases $\frac{g(\rho)}{1-g(\rho)}$ becomes more negative and the price becomes more important for the TV choice.

One Cent Thresholds

An interesting feature of the data is that 18% of the submitted thresholds are one cent below the current price of the TV which is one of the predefined options on camelcamelcamel.com. This mass point is not predicted by the model laid out so far. To incorporate these thresholds into estimation they are interpreted as a corner solution which is chosen by consumers with valuations $v_{ij} > \frac{P_{jt}}{s(\rho)}$. The associated expected utility is assumed to be equal to the utility of purchasing immediately $v_{ij} - P_{jt}$. An interpretation of this assumption is that users of the price alert service no longer consider the possibility of buying immediately and therefore users with high valuations submit the highest possible threshold. A concern with this interpretation is that it leads us to overestimate the users' valuations. In section 6.5 I conduct a robustness check and estimate consumer preferences without using thresholds one cent below the current price.

Discussion

The model does not apply to consumers who submit thresholds for multiple TVs or to consumers who submit multiple thresholds for one TV which is sold by Amazon and a 3rd party retailer. In the data 59% of the users of camelcamelcamel.com submit a threshold for a single TV and a single retailer. In the model consumer preferences are stable over time and the optimal stopping problem is solved by a threshold rule. Therefore the model cannot explain why some users of camelcamelcamel.com update their initial thresholds. This restriction excludes another 4% of the users from the analysis.

Empirically the most relevant of these limitations is that the model does not apply to consumers who submit thresholds for multiple TVs. Most of these consumers submit thresholds for two (18%) or three (7%) TVs. The optimal stopping problem of a consumer

who constantly monitors the prices of two or three TVs is tractable. However, the optimal stopping policy for this problem is not described by a set of threshold prices because the utility of waiting depends on the prices of all tracked TVs. This raises the question how to interpret the submitted price alert thresholds in this case. One possible interpretation is that the consumer restricts herself to policies of the threshold type and plans to purchase the first TV that reaches its threshold price. When the consumer arrives at the market she must compare all TV pairs or triples to decide which TVs to select. This results in a very large choice set. For example in this paper the consumer chooses between 52 TVs which results in 1326 TV pairs and 22100 TV triples. The size of the choice set makes it computationally burdensome to include these consumers into estimation.

I assume that the consumer commits to purchase the TV for which she submitted the threshold and no longer considers buying one of the other TVs. A natural concern is that the consumer reoptimizes when she receives the price alert. The consumers could check the prices of the other TVs again when the price threshold is reached and purchase a different TV if it is a better choice at the new prices. As I do not observe purchases this assumption cannot be evaluated directly. Instead I evaluate the plausibility of this assumption indirectly after the model has been estimated: In section 6.3 I find that only few consumers would prefer to purchase a different TV if they check the prices of the other TVs again when their threshold price has been reached. This suggests that it is plausible to assume that the consumer commits to the TV for which she submitted the threshold.

4 Identification of Preferences

This section discusses identification of the distribution of valuations for different TVs. As in much of the literature estimating dynamic models, the discount rate is assumed to be known by the econometrician. Identification of the functions $\bar{P}(\cdot, \rho)$ and $W(\cdot, \cdot, \rho)$ is not discussed because they are expectations over the price process and therefore directly identified from price data.

The time when a consumer enters the market and submits a threshold is assumed to be exogenous. Therefore consumers face exogenous price variation.

The primitive of interest is the density of valuations f on the support $[\underline{v}_1, \bar{v}_1] \times \dots \times [\underline{v}_J, \bar{v}_J]$ with $\underline{v}_1, \dots, \underline{v}_J \geq 0$. The identification problem is similar to the identification of the Roy Model in Heckman and Honore (1990). The identification argument exploits the fact that $\bar{P}(v_{ij}, \rho)$ is increasing in v_{ij} and $W(v_{ij}, P_{jt}, \rho)$ is increasing in v_{ij} and decreasing in P_{jt} . The proof does however not rely on the particular functional forms of \bar{P} and W under the price process defined in Assumption 1.

Let $\nu(v_{1i}, P_{1t}, P_{jt})$ be the valuation for TV j such that consumer i is indifferent between TVs 1 and j , i.e. $W(\nu(v_{1i}, P_{1t}, P_{jt}), P_{jt}, \rho) = W(v_{1i}, P_{1t}, \rho)$. Identification of ν follows

immediately from identification of W and its monotonicity in the first argument.

Consider the probability that a consumer chooses TV 1 and submits a price threshold smaller than $p < P_{1t}$ for a particular vector of prices. This probability is directly identified from the consumer choice data for different values of p and some region over which we observe prices to vary.

$$\begin{aligned} & \Pr\left(\overline{P}(v_{i1}, \rho) \leq p, \text{ TV 1 chosen} | P_{1t}, \dots, P_{Jt}\right) \\ &= \Pr\left(v_{i1} \leq \frac{p}{s(\rho)}, v_{ij} \leq \nu(v_{1i}, P_{1t}, P_{jt}) \forall j\right) \\ &= \int_{\underline{v}_1}^{\frac{p}{s(\rho)}} \int_{\underline{v}_2}^{\nu(u_1, P_{1t}, P_{2t})} \dots \int_{\underline{v}_J}^{\nu(u_1, P_{1t}, P_{Jt})} f(u_1, u_2, \dots, u_J) du_J \dots du_2 du_1 \end{aligned}$$

Take the derivative with respect to p to obtain

$$\frac{1}{s(\rho)} \int_{\underline{v}_2}^{\nu\left(\frac{p}{s(\rho)}, P_{1t}, P_{2t}\right)} \dots \int_{\underline{v}_J}^{\nu\left(\frac{p}{s(\rho)}, P_{1t}, P_{Jt}\right)} f\left(\frac{p}{s(\rho)}, u_2, \dots, u_J\right) du_J \dots du_2$$

Now take derivatives with respect to P_{jt} for $j = 2 \dots J$ which yields

$$\frac{1}{s(\rho)} \frac{\partial \nu\left(\frac{p}{s(\rho)}, P_{1t}, P_{2t}\right)}{\partial P_{2t}} \dots \frac{\partial \nu\left(\frac{p}{s(\rho)}, P_{1t}, P_{Jt}\right)}{\partial P_{Jt}} f\left(\frac{p}{s(\rho)}, \nu\left(\frac{p}{s(\rho)}, P_{1t}, P_{2t}\right), \dots, \nu\left(\frac{p}{s(\rho)}, P_{1t}, P_{Jt}\right)\right) \quad (4)$$

As s and ν are known functions this reveals the density of the valuations at one particular point.

Using variation in P_{jt} and by varying p we can trace out the density at other points. To identify f at some point v_1, \dots, v_J , choose P_{1t} large enough such that $P_{1t} > v_1 s(\rho)$ and $W(v_1, P_{1t}, \rho) < v_j$ for $j = 2 \dots J$. Let $p = v_1 s(\rho)$ and pick P_{jt} such that $W(v_1, P_{1t}, \rho) = W(v_j, P_{jt}, \rho)$ for $j = 2 \dots J$.

Proposition 1 (Identification). *The joint distribution of valuations is nonparametrically identified on a region where price variation is observed.*

Remark 2. The mapping from the region on which price variation is observed to the region on which the density is identified is clear from (4). Sufficient price variation to trace out the density on its whole support is a strong requirement. In most applications to differentiated product markets the dimension of the valuation density is large. For instance, in this paper there are $J = 52$ TVs in the choice set. For estimation I therefore specify valuations as a function of product characteristics with random coefficients. Intuitively the distribution of random coefficients is identified if the joint distribution of valuations is identified because the space of product characteristics has a lower dimension than J and therefore there is

a unique distribution of random coefficients which generates a particular distribution of valuations.

5 Estimation

Estimation proceeds in two steps. In the first step data on price histories is used to estimate s and \widetilde{W} . In the second step consumer choice data is used to estimate preferences for TV characteristics.

I focus on the six months between 8/1/2011 and 1/31/2012. To keep the size of the choice set manageable I focus on 3D TVs produced by one of the four large brands Samsung, LG, Panasonic and Sony which are sold by Amazon directly. These restrictions leave a subsample of 706 consumers who choose from 52 TVs. Table 3 in Appendix D contains some summary statistics of TV prices during the sample period.

5.1 First Step: Estimate s and \widetilde{W}

To estimate s and \widetilde{W} I use data on the price histories of the 52 TVs in the choice set. Both, s and \widetilde{W} are expectations over

$$Z = \exp \left(\inf_{0 \leq t' \leq \tau(\rho)} X_{jt'} \right),$$

where $\tau(\rho)$ is an exponential random variable with mean $\frac{1}{\rho}$ and X_{jt} is a compound Poisson process with arrival rate λ and jump size distribution G . The expressions for s and \widetilde{W} are given in equations (1) and (2). The annual discount factor is assumed to be 0.9 which corresponds to a daily discount rate of $\rho = 2.89e-4$.¹⁴ A sensitivity analysis with alternative discount rates is conducted in section 6.4.

To obtain estimates of $s(\rho)$ and $\widetilde{W}(\cdot, \rho)$ we generate many draws of Z and replace the expectation operator with the average simulation outcome. The average number of price changes per day is 1.96 which corresponds to $\hat{\lambda} = 0.509$. A smooth bootstrap procedure is used to generate price changes for the simulation of Z : Price changes are drawn with replacement from the price data and a small normally distributed smoothing term is added to every draw to obtain a continuous distribution of price changes. The degree of smoothing is controlled by a bandwidth parameter set to $h = 1.06\sigma N^{-0.2}$, where σ is the sample standard deviation of G and N is the sample size of price changes. The result is robust to changes in the smoothing parameter. A conventional bootstrap without smoothing ($h = 0$) produces similar results.

¹⁴This value is chosen to be similar to Gowrisankaran and Rysman (2012) who consider a monthly discount factor of 0.99.

5.2 Second Step: Estimate Preferences

Valuations are specified as a function of TV characteristics:

$$v_{ij} = \exp(\beta_i x_j + \epsilon_{ij}),$$

where x_j are the characteristics of TV j and ϵ_{ij} is a normal error which is iid across consumers and TVs with standard deviation σ_ϵ . This specification ensures that valuations are positive and therefore each consumer can obtain nonnegative expected utility by submitting a positive price alert threshold.

Each brand offers various series of TVs aimed at different market segments. Panasonic for example, offers the ST30 entry series, the GT30 midrange series and the VT30 high end series. Each series is offered with different screen sizes. To capture these differences the vector of product characteristics x_j contains dummies for brands, the series and the screen size. The coefficients on brand dummies and screen size are normally distributed while the series dummies are associated with nonrandom coefficients. As Samsung offers most of their series either in a LCD or a Plasma version, a Samsung-Plasma dummy is also included.

The estimation method is maximum likelihood. Observations where the user chose a corner solution, i.e. the price threshold is 1 cent below the current price and observations where the user chose an interior solution must be treated differently.

For interior solutions we can use the estimated optimal threshold function \bar{P} to invert the observed price alert threshold and obtain the underlying valuation for the chosen TV. Using the estimate of W we can also recover the expected utility the user obtains from submitting this threshold. Lastly, we can obtain upper bounds on valuations for all TVs the user did not choose.

For corner solutions we only obtain a lower bound on the valuation of the chosen TV and we have to integrate out over possible valuations for the chosen TV.

The likelihood function for interior and corner solutions is given in Appendix B. The unobservables are integrated out with a quadrature method. The KNITRO solver is used to maximize the likelihood function. Different starting values which I tried all lead to the same estimates.

The estimates are presented in Table 4 in Appendix D. At the mean of the screensize coefficient 1 inch of screensize leads to an increase in valuation of 0.64%. Two of the brand dummies are associated with nonrandom coefficients because the estimate of the standard deviation goes to zero with random coefficients. The standard deviation on the Samsung and the Sony dummy correspond to a 28% and 46% increase in valuation, respectively. The model dummies have the expected order except for the Panasonic GT30 which is estimated to be less valuable than Panasonic's entry model the ST30. The difference between these two models, however, is not statistically significant. The standard deviation of the error

term corresponds to a 28% increase in valuations.

6 Implications of the Estimates

6.1 Substitution

If the price of a TV is increased the utility from submitting a threshold for the TV decreases but the optimal threshold for the TV does not change. Therefore consumers respond to the price increase either by submitting a threshold for a different TV or by delaying their purchase of the same TV to a later time.

Table 1 summarizes the substitution patterns across TVs and over time in response to a price increase for a particular TV. The results are obtained by simulating the arrival of a large number of consumers at every day in the sample period and by selectively increasing the price of a particular model by 3%.¹⁵ The price increase is permanent in the sense that all possible realizations of future prices are also increased. However, as prices tend to fall over time, prices will eventually return to the level without the price increase. In the baseline case the arrival rate of consumers is assumed to be uniform throughout the sample period. Table 13 in Appendix D shows results if the arrival rate is estimated which differ only slightly from the results with a uniform arrival rate.

Table 1 shows the flow of consumers who would have bought a particular TV within one month after their arrival but no longer do so if its price is increased.¹⁶ Consumers who chose a corner solution are not considered.¹⁷ Column (4) is the fraction of consumers who buy the same TV after the first month and column (5) are the remaining consumers who chose a different TV.¹⁸

The results show that the possibility to delay the purchase is important. Most of the substitution (71%-89%) is towards the same TV at a later time. TVs compete mostly with their own future selves rather than other TVs. Substitution towards other TVs is more important for entry models such as the 32' Sony EX720 and less important for high-end TVs such as the 65' Samsung D8000.

These substitution patterns differ from Conlon (2010) who estimates dynamic demand for LCD TVs using a quarterly panel of prices and market shares in a dynamic discrete

¹⁵The average absolute price change during the sample period is 2.7%. The results are similar if the price is increased by 1% or by 10% (not reported). As the price increase becomes larger substitution towards other TVs becomes somewhat more important.

¹⁶Tables 5 and 6 show results if the time window is decreased to one week or increased to ten weeks.

¹⁷For these consumers the the importance of substitution towards a later purchase time of the same TV depends much more on the magnitude of the price change. If the price increase is small such that they would still submit a corner solution at the new price they do not delay their purchase in response to the price increase. Only for large price increases their optimal threshold is lower than the new price such that they delay their purchase compared to the original price.

¹⁸The time when the threshold is reached depends on the realization of the price process. The results are obtained by integrating out possible realizations of the price path through simulation.

choice framework. Conlon reports substitution in response to a transitory price increase for one quarter. Even though Conlon considers a transitory price change he finds substitution towards the same model at a later time to be relatively unimportant. He points out that this is likely related to the logit error in the dynamic discrete choice framework which is redrawn every period. If preferences for TVs change over time it is less attractive to delay the purchase of a product.

6.2 Buy Now Counterfactual

This section presents results from a counterfactual Buy Now scenario in which intertemporal substitution is shut down and consumers buy a TV immediately when they arrive at the market. In the Buy Now scenario consumers can choose an outside option to guarantee that they get nonnegative utility. The market share of the outside option however is very close to zero because the estimated valuations are large enough to ensure that almost all consumers can find a TV which costs less than their valuation.

The Buy Now counterfactual is a useful benchmark to evaluate how much consumers gain from delaying their purchase. Average consumer surplus increases by \$104 or 1.7% if consumers submit a threshold for a single TV rather than buying immediately upon arrival at the market.¹⁹

To evaluate how much of the potential gains from delaying the purchase are realized by tracking the price of a single TV we would like to consider a consumer who constantly monitors the prices of all TVs and waits for the optimal time to buy. Solving this optimal stopping problem is not feasible without simplifying assumptions on beliefs because the dimension of the state space is too large if the consumer keeps track of all prices. One assumption that is frequently used in the dynamic demand literature to reduce the dimension of the state space is perfect foresight (e.g. Gowrisankaran and Rysman (2012) and Conlon (2010) for durable goods or Hendel and Nevo (forthcoming) for storable goods). Here perfect foresight under different decision horizons is used as a benchmark to put the gain of \$104 into perspective. The decision horizon restricts the number of days over which is optimized. For example if the decision horizon is 1 week the consumer considers the prices of all TVs within 1 week of her arrival at the market. Figure 5 shows that a gain of \$104 can be achieved with a decision horizon of 23 days. Notice however that the increase in average consumer surplus is a concave function of the decision horizon because consumers with high valuations prefer to buy early and do not benefit from a longer horizon. Extending the decision horizon to 120 days, for example, improves the gain to \$174. As gains under

¹⁹Consumer surplus is obtained by simulating the arrival of a large number of consumers on every day in the sample period. In the baseline case the arrival rate of consumers is assumed to be uniform throughout the sample period. The surplus of consumer i arriving at time t is $\arg \max_j v_{ij} - P_{jt}$ in the Buy Now scenario, and $\arg \max_j W(v_{ij}, P_{jt}, \rho)$ if they submit a threshold and delay their purchase.

perfect foresight are an upper bound on gains under more realistic beliefs this suggests that consumers cannot benefit much from tracking the prices of multiple TVs. This could explain why most users of the price alert service submit a threshold for a single TV.




Figure 5: Increase in consumer surplus for the average user under perfect foresight compared to the Buy Now scenario. The decision horizon restricts the number of days after arrival at the market over which is optimized. The graph shows an average over the sample period where each day receives equal weight. The dashed lines mark the increase in consumer surplus if the consumer tracks the price of a single TV with the help of the price alert service (\$104) which corresponds to a decision horizon of 23 days under perfect foresight.

The Buy Now counterfactual is also helpful to understand how intertemporal substitution affects competition between TVs.

Table 2 illustrates that the Buy Now counterfactual predicts own-price elasticities and market shares which differ considerably from the model with intertemporal substitution. Own-price elasticities with respect to a permanent price increase and market shares are obtained through the same simulation procedure described in the previous section.

First consider the predicted own-price elasticities i.e. how many consumers are lost to other TVs in response to a permanent price increase. Column (6) in Table 2 shows that this elasticity is predicted to be higher for all TVs if we shut down intertemporal substitution. The deviation from the model with intertemporal substitution is larger for expensive high end TVs for which most of the substitution is towards the same TV at a later time. For the 65" Samsung-LED/LCD D8000 and the 65" Panasonic VT30 for example the predicted elasticity is more than twice as large as in the model with intertemporal substitution.

Next consider the predicted market shares. The following proposition tells us that the Buy Now scenario predicts lower market shares for expensive TVs.

Proposition 2. *Consider two TVs with prices $P_{1t} \geq P_{2t}$. If the consumer prefers TV 1 over TV 2 in the Buy Now scenario she also prefers to submit a price threshold for TV 1.*

The proof can be found in Appendix C. The intuition for the result is that expensive models exhibit larger absolute price jumps so delaying the purchase is more valuable for these TVs.

Column (7) in Table 2 shows how much the predicted market shares in the Buy Now scenario deviate from the model with intertemporal substitution. The market shares of high end TVs are predicted too low by up to 49% for the 65" Samsung-LED/LCD D8000. The market shares of entry models are predicted too high by up to 44% for the 32" Sony EX720.

6.3 Incentive to Reoptimize

The subsample of consumers which were selected for estimation decided to submit a threshold for a single TV. I assume that they no longer consider the other available TVs after they submitted the price threshold. A possible concern is that consumers reoptimize when their price threshold is reached. The consumers could check the prices of the other TVs again and purchase a different TV if it is a better choice at the new prices. I find that only 4.5% of the users would prefer a different TV if they check the prices of all other TVs once their threshold is reached. Most of the consumers switch to one of three TVs which were not available when they submitted their threshold. If we restrict attention to the TVs which were available when the consumers entered the market only 1% of the consumers would prefer a different TV at the time when the threshold is reached. As only few consumers would benefit from checking the prices of the other TVs again it seems reasonable to assume that the consumers commit to the TV they track.

6.4 Sensitivity Analysis: Alternative Discount Rates

Tables 7 and 8 in Appendix D show estimates of the preference parameters for alternative discount rates. The daily discount rates are $\rho = 3.66e - 4$ and $\rho = 2.14e - 4$ which correspond to annual discount factors of 0.875 and 0.925. While most of the parameters are insensitive to changes in the discount rate the means of the product dummies become larger as the discount rate gets smaller. This is not surprising as the price threshold $\bar{P}(v_{ij}, \rho) = v_{ij}s(\rho)$ is increasing in v_{ij} and ρ . If the discount rate is increased smaller valuations are needed to explain the observed price thresholds.

Tables 9 and 10 in Appendix D show the implications of the estimates for the alternative discount rates. Not surprisingly, a higher discount rate predicts substitution over time to be less important. Substitution towards the same TV after the first month is between 2.3 and 6.6 percentage points smaller for $\rho = 3.66e - 4$ and between 2.4 and 5.5 percentage points larger for $\rho = 2.14e - 4$. The estimated own price elasticities are larger for higher discount

rates. For $\rho = 3.66e - 4$ the elasticities are approximately 20% higher and for $\rho = 2.14e - 4$ approximately 20% lower than in the reference case. Predicted market shares and own-price elasticities in the Buy Now counterfactual differ from the model with intertemporal substitution to a similar extent as in the reference case.

The predicted increase in average consumer surplus compared to the Buy Now counterfactual is \$112 (1.3%) for $\rho = 2.14e - 4$ and \$98 (2.0%) for $\rho = 3.66e - 4$, where the relative change is given in parentheses. Notice that the absolute change is decreasing in ρ , whereas the relative change is increasing in ρ . In absolute terms delaying the purchase is estimated to be more valuable if consumers are more patient. As estimated valuations are decreasing in the discount rate, however, the increase amounts to a larger fraction of consumer surplus for larger ρ .

Overall, the results are not very sensitive to the choice of the discount rate.

6.5 Robustness: Dropping One Cent Thresholds

Table 11 in Appendix D shows estimates of the preference parameters if the price thresholds 1 cent below the current price are dropped. Most of the estimates are similar to the baseline specification. An important difference is that the estimate of σ_ϵ decreases from 0.245 to 0.182, i.e. the taste for particular TVs becomes less important.

Table 12 in Appendix D show the implications of the estimates. As tastes for particular TVs are less important there is more substitution across TVs. The fraction of consumers who substitute towards the same TV at a later time is between 2.3 and 5.4 percentage points lower than in the reference case. For the same reason own price elasticities are between 20% and 40% higher than in the reference case. Predicted market shares and own-price elasticities in the Buy Now counterfactual differ from the model with intertemporal substitution to a similar extent as in the reference case. The predicted increase in average consumer surplus compared to the Buy Now counterfactual is \$112.5 or 1.9%.

Overall, dropping the thresholds one cent below the current price has a moderate effect on the results.

Brand	Series	Screen Size	Buy the same TV after the first month (in %)	Buy a different TV (in %)	
Samsung-LED/LCD	D6400	40	74.6	25.4	
		46	75.4	24.6	
		55	77.3	22.7	
		60	79.9	20.1	
	D6500	40	75.9	24.1	
		46	75.8	24.2	
		55	77.8	22.2	
		60	79.9	20.1	
	D7000	46	80.7	19.3	
		55	82	18	
		60	84.1	15.9	
	D8000	46	82.7	17.3	
		55	84.1	15.9	
		60	86.6	13.4	
		65	88.7	11.3	
Samsung-Plasma	D6500	51	72.7	27.3	
		59	75.1	24.9	
	D7000	51	77.7	22.3	
		59	80.2	19.8	
		64	81.7	18.3	
	D8000	51	79.3	20.7	
		59	81.5	18.5	
		64	83.8	16.2	
LG	LW5600	47	79.4	20.6	
		55	81.4	18.6	
	LW6500	47	80.2	19.8	
		55	83	17	
		65	87.1	12.9	
	Panasonic	ST30	42	75.7	24.3
46			76.5	23.5	
50			76.7	23.3	
55			78.7	21.3	
60			80.3	19.7	
65			83.7	16.3	
GT30		50	77.8	22.2	
		55	79.6	20.4	
		60	81.7	18.3	
		65	84	16	
VT30		55	84.2	15.8	
		65	87.4	12.6	
Sony		EX720	32	72.7	27.3
			40	71.3	28.7
			46	73.2	26.8
			55	74.8	25.2
			60	79.1	20.9
		NX720	46	78.1	21.9
	55		79.9	20.1	
	60		83	17	
	HX820	46	81.7	18.3	
		55	83.9	16.1	
	HX929	46	86	14	
		55	87.6	12.4	

Table 1: Column (4) shows the fraction of consumers who respond to the price increase by purchasing the same TV at a later time. Column (5) are the remaining consumers who choose a different TV.

Brand	Series	Screen Size	Predicted Own-Price Elasticities with intertemporal substitution	Predicted Market Shares with intertemporal substitution (in %)	Buy Now: Elasticity Deviation (in %)	Buy Now: Market Share Deviation (in %)	
Samsung-LED/LCD	D6400	40	-1.42	0.8	10.6	34.7	
		46	-1.49	0.9	16.1	26.4	
		55	-1.63	0.8	30.9	7.2	
		60	-1.70	0.7	46.0	-10.9	
	D6500	40	-1.51	0.7	15.3	29.2	
		46	-1.49	1.0	18.6	23.8	
		55	-1.58	0.8	32.4	7.1	
		60	-1.68	0.6	48.1	-11.8	
	D7000	46	-1.34	2.7	30.2	10.1	
		55	-1.37	2.9	39.4	-2.1	
		60	-1.39	2.6	59.5	-17.3	
	D8000	46	-1.28	3.7	39.7	1.7	
		55	-1.28	4.2	51.6	-8.1	
		60	-1.34	3.0	77.8	-29.5	
		65	-1.36	2.9	102.4	-49.1	
	Samsung-Plasma	D6500	51	-1.60	0.4	14.6	27.5
59			-1.74	0.4	27.2	8.4	
D7000		51	-1.48	1.0	25.9	13.4	
		59	-1.55	1.0	42.4	-5.4	
		64	-1.54	1.3	50.7	-15.8	
D8000		51	-1.39	1.4	26.9	8.7	
		59	-1.43	1.7	42.3	-5.0	
		64	-1.46	1.5	59.4	-20.7	
LG	LW5600	47	-1.15	5.2	12.9	21.2	
		55	-1.23	5.1	27.3	10.3	
	LW6500	47	-1.18	3.4	16.8	19.7	
		55	-1.31	2.7	45.3	-0.5	
		65	-1.30	4.4	78.4	-26.4	
	Panasonic	ST30	42	-1.08	3.5	0.8	32.7
46			-1.18	3.3	6.2	28.7	
50			-1.21	3.1	9.0	24.8	
55			-1.35	3.0	21.4	15.7	
60			-1.37	3.1	31.3	6.4	
65			-1.43	2.5	55.0	-13.7	
GT30		50	-1.44	1.9	19.9	21.0	
		55	-1.50	1.8	35.0	8.3	
		60	-1.57	1.5	55.2	-11.3	
		65	-1.57	1.5	70.7	-27.0	
VT30		55	-1.45	1.8	67.7	-18.4	
		65	-1.40	2.9	102.5	-44.6	
Sony		EX720	32	-1.18	0.3	5.9	43.8
			40	-1.28	0.4	9.0	35.2
	46		-1.42	0.3	13.1	25.8	
	55		-1.52	0.3	22.7	11.1	
	60		-1.67	0.2	41.5	-10.0	
	NX720	46	-1.31	0.9	22.3	15.0	
		55	-1.38	0.9	32.7	1.6	
		60	-1.47	0.7	52.2	-17.2	
	HX820	46	-1.19	1.8	30.4	6.8	
		55	-1.23	1.9	44.1	-6.8	
	HX929	46	-1.21	2.1	66.0	-16.1	
		55	-1.16	2.6	70.7	-21.5	

Table 2: Columns (4) and (5) are own-price elasticities and market shares predicted by the model with intertemporal substitution. Column (6) shows how much the predicted own-price elasticities in the Buy Now counterfactual differ from the model with intertemporal substitution. Column (7) shows how much predicted market shares in the Buy Now scenario deviate from the model with intertemporal substitution.

7 Summary and Conclusion

Consumers planning to purchase durable goods often face volatile and decreasing prices. This creates an incentive for them to delay their purchase and wait for better prices. In this paper I argue that data from a price alert service provides a window into consumer behavior in this environment. I demonstrate how this data can be used to estimate consumer preferences and to quantify how consumers substitute across products and over time.

I propose a simple discrete/continuous choice model in which the price alert threshold is the solution to an optimal stopping problem. Price variation across consumers who enter the market at different times is shown to nonparametrically identify the joint distribution of valuations for different TVs. A two-step estimation procedure for preferences is proposed. The estimation results are used to quantify substitution across TVs and over time. The substitution patterns suggest that products compete to a large extent with their own future selves rather than with other products.

For tractability reasons this paper focuses on the majority of consumers who track the price of a single TV. Intuitively, this restriction means that substitution across TVs is underestimated. Finding a tractable way to incorporate consumers who submit thresholds for multiple TVs is therefore an important task for future work.

Another avenue for future research is to combine data from a price alert service and aggregate purchase data. While aggregate quantity data is not published for products sold on Amazon.com, sales rank data is publically available.

A Approximation of \widetilde{W}

The functional form used for the approximation implies that:

$$\log \widetilde{W} \left(\frac{P_{tj}}{\overline{P}(v_{ij}, \rho)}, \rho \right) \approx g(\rho) \log \left(\frac{P_{tj}}{\overline{P}(v_{ij}, \rho)} \right)$$

To approximate \widetilde{W} , $g(\rho)$ is chosen to match the average slope of the estimated $\log(\widetilde{W})$ on the interval $[1, 4]$:

$$g(\rho) = \frac{\log \widetilde{W}(4, \rho) - \log \widetilde{W}(1, \rho)}{\log(4) - \log(1)}$$

Figure 6 shows the simulated \widetilde{W} in blue and the approximation in red for different values of the discount factor.




Figure 6: Approximation of \widetilde{W} for different discount rates. The estimated \widetilde{W} is shown in blue and the approximation in red. The annual discount factor varies from 0.99 for the highest line to 0.1 for the lowest line. The approximation is better for higher discount factors. For an annual discount factor of 0.9 and $4\overline{P}(v_{ij}, \rho) > P_{tj}$ the maximum approximation error is less than 0.02%. For an annual discount factor of 0.1 the approximation error is still smaller than 0.75%.

B Second Step: Likelihood

Let F_ϵ and f_ϵ be the cdf and the pdf of the error and f_β the pdf of the random coefficients.

For an observation with an interior solution the likelihood takes the following form:

$$\int f_{\beta}(\beta_i) f_{\epsilon}(\epsilon^*(\bar{P}_{ij}, \beta_i)) \times \frac{\partial}{\partial \bar{P}} \epsilon^*(\bar{P}_{ij}, \beta_i) \times \prod_{k \neq j} F_{\epsilon}(\bar{\epsilon}(\bar{P}_{ij}, \beta_i, P_{jt}, P_{kt})) d\beta_i,$$

where \bar{P}_{ij} is the threshold for the chosen TV, $\epsilon^*(\bar{P}_{ij}, \beta_i)$ is the error for TV j which generates \bar{P}_{ij} for β_i and $\bar{\epsilon}(\bar{P}_{ij}, \beta_i, P_{jt}, P_{kt})$ are the errors for $k \neq j$ such that the consumer would be indifferent between j and k .

For an observation with a corner solution the error of the chosen TV must also be integrated out:

$$\int \int_{\epsilon^*(P_{j0}, \beta_i)}^{\infty} f_{\beta}(\beta_i) f_{\epsilon}(\epsilon) \times \prod_{k \neq j} F_{\epsilon}(\bar{\epsilon}(\epsilon, \beta_i, P_{jt}, P_{kt})) d\epsilon d\beta_i,$$

where $\epsilon^*(P_{jt}, \beta_i)$ is the lowest error generating a corner solution and $\bar{\epsilon}(\epsilon, \beta_i, P_{jt}, P_{kt})$ are the errors for $k \neq j$ such that the consumer would be indifferent between j and k .

C Proof of Proposition 2

Proof. As the consumer prefers TV 1 in the Buy Now scenario

$$V_1 - P_{1t} \geq V_2 - P_{2t} \tag{5}$$

First consider the case where $s(\rho) V_1 \leq P_{1t}$ and $s(\rho) V_2 \leq P_{2t}$:

$$\begin{aligned} & s(\rho) [W(V_1, P_{1t}, \rho) - W(V_2, P_{2t}, \rho)] \\ &= \mathbb{E}[\max\{V_1 s(\rho) - P_{1t} Z, 0\}] - \mathbb{E}[\max\{V_2 s(\rho) - P_{2t} Z, 0\}] \\ &\geq \mathbb{E}[\max\{s(\rho)(V_2 + P_{1t} - P_{2t}) - P_{1t} Z, 0\}] - \mathbb{E}[\max\{V_2 s(\rho) - P_{2t} Z, 0\}] \\ &\geq 0, \end{aligned}$$

where the first inequality follows from (5). The second inequality can be shown pointwise in Z : If $Z \geq s(\rho)$ then $V_2 s(\rho) - P_{2t} Z \leq 0$. If $Z \leq s(\rho)$, then $s(\rho)(V_2 + P_{1t} - P_{2t}) - P_{1t} Z \geq V_2 s(\rho) - P_{2t} Z$.

Second, if $s(\rho) V_1 \leq P_{1t}$ and $s(\rho) V_2 > P_{2t}$ then the utility from TV 2 is $V_2 - P_{2t}$ and $W(V_1, P_{1t}, \rho) \geq V_1 - P_{1t} \geq V_2 - P_{2t}$.

Third, if $s(\rho) V_1 > P_{1t}$ and $s(\rho) V_2 \leq P_{2t}$. This implies that $V_1 - P_{1t} > \frac{P_{1t}(1-s(\rho))}{s(\rho)}$ and at the same time $W(V_2, P_{2t}, \rho) \leq V_2(1-s(\rho)) \leq \frac{P_{2t}(1-s(\rho))}{s(\rho)}$. It follows from $P_{1t} \geq P_{2t}$ that the consumer prefers TV 1.

Lastly, the case where $s(\rho) V_1 > P_{1t}$ and $s(\rho) V_2 > P_{2t}$ is identical to the Buy Now

scenario, which concludes the proof. □

D Additional Figures and Tables




Figure 7: Screenshot of a product page on camelcamelcamel.com. The user sees the price history and can submit her price threshold which is called the desired price. Price alerts are either delivered via email or Twitter.




Figure 8: Distribution of the gap between the price and the submitted threshold (in \$) with a mass point at 1 cent. The last bar collects all values which are larger than \$500.

Brand	Series	Screen Size	Minimum Price (in \$)	Maximum Price (in \$)	Average Daily Price Correlation	
Samsung-LED/LCD	D6400	40	730	1018	0.12	
		46	940	1250	0.07	
		55	1449	1800	0.05	
		60	1852	2428	0.07	
	D6500	40	952	1200	0.01	
		46	900	1800	0.02	
		55	1397	1900	0.05	
		60	2099	2454	0.00	
	D7000	46	1168	1719	0.10	
		55	1639	2789	0.02	
		60	2116	2669	0.01	
	D8000	46	1293	2520	0.08	
55		1822	3149	0.05		
60		2582	3869	0.04		
65		3299	5400	0.04		
Samsung-Plasma	D6500	51	900	1440	0.01	
		59	1328	1956	0.07	
	D7000	51	1032	1900	0.06	
		59	1656	2249	0.07	
		64	2014	3016	0.07	
	D8000	51	1231	2300	0.07	
		59	1800	2162	0.09	
		64	2269	2835	0.11	
	LG	LW5600	47	898	1120	0.01
55			1268	1499	0.00	
LW6500		47	1040	1300	-0.01	
		55	1425	2100	-0.02	
		65	2299	3499	-0.01	
Panasonic		ST30	42	660	800	0.04
	46		700	985	0.08	
	50		805	1139	0.10	
	55		1100	1384	0.09	
	60		1250	1759	0.10	
	65		1782	2382	0.07	
	GT30	50	975	1447	0.09	
		55	1300	1783	0.07	
		60	1640	2800	0.07	
		65	2320	2700	0.05	
	VT30	55	1775	2562	0.08	
		65	2900	4300	0.02	
	Sony	EX720	32	695	754	0.03
			40	739	1168	0.05
			46	960	1146	0.01
			55	1260	1667	0.04
			60	1939	2788	0.04
		NX720	46	924	2100	0.06
55			1500	2900	0.05	
60			2299	3500	0.04	
HX820		46	1299	2600	0.05	
		55	1700	2798	0.05	
HX929		46	2000	2531	0.01	
		55	2398	3049	-0.10	

Table 3: Summary statistics for the prices of the 52 TVs in the choice set during the sample period. Columns (4) and (5) show the minimum and the maximum price of the TV during the sample period. Column (6) shows the average correlation of daily price changes. The price changes are measured in % and are based on the minimum price for each day. These price changes are used to compute a matrix of correlation coefficients. Column (6) shows the average correlation coefficient for each TV.

			Coefficient	Std. Err.
Screen Size	Screen Size	mean	0.0064	0.0008
		std. dev.	0.0049	0.0002
Brand Dummies	Samsung	mean	7.8370	0.0580
		std. dev.	0.2476	0.0208
	LG	constant	8.0513	0.0479
	Panasonic	constant	7.9696	0.0488
	Sony	mean	7.4851	0.0955
std. dev.		0.3835	0.0548	
Series Dummies	Samsung D6500	constant	0.0190	0.0305
	Samsung D7000	constant	0.1659	0.0274
	Samsung D8000	constant	0.2321	0.0269
	Samsung Plasma	constant	-0.1095	0.0172
	LG 6500	constant	0.0217	0.0236
	Panasonic GT30	constant	-0.0311	0.0198
	Panasonic VT30	constant	0.0917	0.0252
	Sony NX720	constant	0.1402	0.0441
	Sony HX820	constant	0.2633	0.0435
	Sony HX929	constant	0.3490	0.0437
ϵ Std. Dev.	σ_ϵ	constant	0.2456	0.0146

Table 4: Estimates of the preference parameters and asymptotic standard errors for the baseline specification with $\rho = 2.89e - 4$. The left out series dummies are the entry models of each brand: Samsung D6400, LG 5600, Panasonic ST30 and Sony EX720. The coefficient on the LG and Panasonic brand dummies are constant because the variance estimate is zero if the model is estimated with random coefficients.

Brand	Series	Screen Size	Buy the same TV after the first week (in %)	Buy a different TV (in %)	
Samsung-LED/LCD	D6400	40	87.4	12.6	
		46	87.8	12.2	
		55	88.7	11.3	
		60	90.2	9.8	
	D6500	40	87.9	12.1	
		46	87.8	12.2	
		55	89.1	10.9	
		60	90.1	9.9	
	D7000	46	90.7	9.3	
		55	91.4	8.6	
		60	92.4	7.6	
	D8000	46	91.7	8.3	
55		92.5	7.5		
60		93.8	6.2		
65		95.1	4.9		
Samsung-Plasma	D6500	51	86	14	
		59	87.4	12.6	
	D7000	51	88.9	11.1	
		59	90.4	9.6	
		64	91.2	8.8	
	D8000	51	89.9	10.1	
		59	91.1	8.9	
		64	92.3	7.7	
	LG	LW5600	47	90.1	9.9
55			91.1	8.9	
LW6500		47	90.5	9.5	
		55	91.7	8.3	
		65	94.1	5.9	
Panasonic	ST30	42	88.2	11.8	
		46	88.5	11.5	
		50	88.7	11.3	
		55	89.6	10.4	
		60	90.5	9.5	
		65	92.3	7.7	
	GT30	50	89.1	10.9	
		55	90	10	
		60	91	9	
		65	92.4	7.6	
	VT30	55	92.5	7.5	
		65	94.4	5.6	
	Sony	EX720	32	86.7	13.3
			40	85.4	14.6
			46	86.6	13.4
55			87.3	12.7	
60			90	10	
NX720		46	89.3	10.7	
		55	90.3	9.7	
		60	91.9	8.1	
HX820		46	91.2	8.8	
		55	92.4	7.6	
HX929		46	93.4	6.6	
		55	94.4	5.6	

Table 5: The flows of consumers who in absence of a price increase would have bought the TV within one week of their arrival.

Brand	Series	Screen Size	Buy the same TV after the first ten weeks (in %)	Buy a different TV (in %)	
Samsung-LED/LCD	D6400	40	60	40	
		46	61.1	38.9	
		55	63.9	36.1	
		60	67.6	32.4	
	D6500	40	61.9	38.1	
		46	61.8	38.2	
		55	64.4	35.6	
		60	67.7	32.3	
	D7000	46	68.4	31.6	
		55	70.3	29.7	
		60	73.4	26.6	
	D8000	46	71.4	28.6	
55		73.4	26.6		
60		76.8	23.2		
65		79	21		
Samsung-Plasma	D6500	51	57.9	42.1	
		59	61	39	
	D7000	51	64.3	35.7	
		59	67.9	32.1	
		64	70	30	
	D8000	51	66.4	33.6	
		59	69.5	30.5	
		64	72.8	27.2	
	LG	LW5600	47	66.3	33.7
55			69.1	30.9	
LW6500		47	67.4	32.6	
		55	71.8	28.2	
		65	77.5	22.5	
Panasonic	ST30	42	61.1	38.9	
		46	62.2	37.8	
		50	62.6	37.4	
		55	65.4	34.6	
		60	67.8	32.2	
		65	72.6	27.4	
	GT30	50	64.4	35.6	
		55	67	33	
		60	70.2	29.8	
		65	73.2	26.8	
	VT30	55	73.6	26.4	
		65	77.6	22.4	
	Sony	EX720	32	57.1	42.9
			40	55.7	44.3
			46	58.2	41.8
55			60.7	39.3	
60			66.3	33.7	
NX720		46	64.6	35.4	
		55	67.2	32.8	
		60	71.7	28.3	
HX820		46	69.7	30.3	
		55	72.9	27.1	
HX929		46	76.3	23.7	
		55	78.2	21.8	

Table 6: The flows of consumers who in absence of a price increase would have bought the TV within ten weeks of their arrival.

			Coefficient	Std. Err.
Screen Size	Screen Size	mean	0.0079	0.0007
		std. dev.	0.0047	0.0002
Brand Dummies	Samsung	mean	7.6131	0.0558
		std. dev.	0.2387	0.0200
	LG	constant	7.8123	0.0462
	Panasonic	constant	7.7280	0.0469
	Sony	mean	7.2713	0.0908
		std. dev.	0.3692	0.0513
Series Dummies	Samsung D6500	constant	0.0193	0.0294
	Samsung D7000	constant	0.1661	0.0263
	Samsung D8000	constant	0.2350	0.0260
	Samsung Plasma	constant	-0.1156	0.0166
	LG 6500	constant	0.0257	0.0227
	Panasonic GT30	constant	-0.0227	0.0190
	Panasonic VT30	constant	0.1077	0.0242
	Sony NX720	constant	0.1414	0.0425
	Sony HX820	constant	0.2644	0.0420
	Sony HX929	constant	0.3606	0.0422
ϵ Std. Dev.	σ_ϵ	constant	0.2365	0.0142

Table 7: Estimates of the preference parameters for $\rho = 3.66e - 4$.

			Coefficient	Std. Err.
Screen Size	Screen Size	mean	0.0048	0.0008
		std. dev.	0.0051	0.0002
Brand Dummies	Samsung	mean	8.1237	0.0597
		std. dev.	0.2567	0.0213
	LG	constant	8.3521	0.0492
	Panasonic	constant	8.2737	0.0501
	Sony	mean	7.7614	0.1000
		std. dev.	0.3979	0.0583
Series Dummies	Samsung D6500	constant	0.0186	0.0314
	Samsung D7000	constant	0.1640	0.0281
	Samsung D8000	constant	0.2269	0.0277
	Samsung Plasma	constant	-0.1020	0.0177
	LG 6500	constant	0.0173	0.0243
	Panasonic GT30	constant	-0.0394	0.0203
	Panasonic VT30	constant	0.0743	0.0259
	Sony NX720	constant	0.1380	0.0453
	Sony HX820	constant	0.2602	0.0447
	Sony HX929	constant	0.3339	0.0450
ϵ Std. Dev.	σ_ϵ	constant	0.2526	0.0148

Table 8: Estimates of the preference parameters for $\rho = 2.14e - 4$.

Brand	Series	Screen Size	Buy the same TV after the first month (in %)	Buy a different TV (in %)	Own-price Elasticity with i.s.	Market Shares with i.s. (in %)	Buy Now: Own-price Elasticity Deviation (in %)	Buy Now: Market Share Deviation (in %)	
Samsung-LED/LCD	D6400	40	70.6	29.4	-1.76	0.8	9.0	41.2	
		46	71.3	28.7	-1.85	0.9	14.0	30.7	
		55	73.8	26.2	-1.98	0.8	30.0	7.5	
		60	75.7	24.3	-2.09	0.7	42.0	-13.0	
	D6500	40	72	28	-1.85	0.7	16.4	34.4	
		46	71.7	28.3	-1.83	1.0	18.0	27.4	
		55	73.8	26.2	-1.93	0.8	28.3	7.2	
		60	76.6	23.4	-2.03	0.6	46.7	-14.3	
	D7000	46	77.1	22.9	-1.65	2.6	27.5	10.9	
		55	78.8	21.2	-1.64	3.0	40.1	-3.0	
		60	81.2	18.8	-1.70	2.6	55.4	-20.3	
	D8000	46	79.5	20.5	-1.57	3.6	38.6	1.2	
55		81.4	18.6	-1.55	4.4	48.4	-10.1		
60		83.9	16.1	-1.63	3.0	71.9	-33.4		
65		86.4	13.6	-1.68	2.8	93.0	-54.0		
Samsung-Plasma	D6500	51	68	32	-1.96	0.4	15.4	31.9	
		59	71.2	28.8	-2.14	0.4	24.5	8.9	
	D7000	51	73.6	26.4	-1.83	1.0	23.3	14.8	
		59	76.5	23.5	-1.93	0.9	37.8	-7.0	
		64	78.4	21.6	-1.88	1.3	46.3	-18.1	
	D8000	51	75.5	24.5	-1.69	1.4	27.0	9.4	
		59	78.2	21.8	-1.75	1.8	38.4	-6.4	
		64	80.8	19.2	-1.78	1.5	55.7	-24.0	
	LG	LW5600	47	75.9	24.1	-1.43	5.1	11.8	24.5
55			78.1	21.9	-1.51	5.2	25.6	11.1	
LW6500		47	76.9	23.1	-1.46	3.5	15.0	22.4	
		55	79.9	20.1	-1.60	2.7	43.2	-1.5	
		65	84.4	15.6	-1.59	4.4	72.7	-30.4	
Panasonic		ST30	42	71.7	28.3	-1.35	3.4	-1.1	39.2
	46		72.7	27.3	-1.47	3.3	4.9	33.9	
	50		72.9	27.1	-1.49	3.2	8.0	28.8	
	55		75	25	-1.65	3.1	21.7	17.6	
	60		76.9	23.1	-1.66	3.1	30.1	6.6	
	65		80.4	19.6	-1.74	2.4	53.6	-16.4	
	GT30	50	73.9	26.1	-1.75	2.0	19.1	23.8	
		55	75.8	24.2	-1.85	1.8	31.9	8.5	
		60	78.3	21.7	-1.91	1.5	52.7	-13.8	
		65	80.9	19.1	-1.92	1.4	66.5	-31.1	
		65	80.9	19.1	-1.92	1.4	66.5	-31.1	
	VT30	55	81	19	-1.76	1.8	65.7	-21.7	
		65	84.9	15.1	-1.71	2.9	93.7	-49.2	
	Sony	EX720	32	66.1	33.9	-1.53	0.2	0.6	54.7
			40	67.5	32.5	-1.55	0.4	9.9	42.5
			46	69.5	30.5	-1.79	0.3	14.3	30.3
			55	71.3	28.7	-1.87	0.3	21.1	12.5
			60	76.2	23.8	-2.05	0.2	37.7	-11.9
NX720		46	74	26	-1.61	0.9	20.3	17.1	
		55	76.4	23.6	-1.66	0.9	31.9	1.4	
		60	79.8	20.2	-1.81	0.7	44.8	-19.9	
HX820		46	78.4	21.6	-1.48	1.8	28.7	7.5	
		55	80.9	19.1	-1.48	1.9	43.7	-8.2	
HX929		46	83.5	16.5	-1.50	2.0	62.8	-18.9	
		55	85.2	14.8	-1.42	2.7	66.4	-24.7	

Table 9: Sensitivity analysis for $\rho = 3.66e - 4$. Columns (4) and (5) show the flows of consumers who in absence of a price increase would have bought the TV within one month of their arrival. Columns (6) and (7) are own-price elasticities and market shares predicted by the model with intertemporal substitution. Columns (8) and (9) show how much predicted own-price elasticities and market shares in the Buy Now counterfactual deviate from the model with intertemporal substitution.

Brand	Series	Screen Size	Buy the same TV after the first month (in %)	Buy a different TV (in %)	Own-price Elasticity with i.s.	Market Shares with i.s. (in %)	Buy Now: Elasticity Deviation (in %)	Buy Now: Market Share Deviation (in %)	
Samsung-LED/LCD	D6400	40	79.3	20.7	-1.09	0.9	12.2	27.7	
		46	79.8	20.2	-1.15	0.9	18.4	21.9	
		55	81.4	18.6	-1.26	0.8	33.3	6.6	
		60	83.4	16.6	-1.30	0.7	52.8	-8.4	
	D6500	40	80.4	19.6	-1.14	0.8	17.8	23.8	
		46	80.3	19.7	-1.16	1.0	20.8	19.8	
		55	81.7	18.3	-1.24	0.7	32.8	6.7	
		60	83.7	16.3	-1.31	0.6	51.2	-9.2	
	D7000	46	84.4	15.6	-1.04	2.7	31.8	8.9	
		55	85.4	14.6	-1.06	2.8	45.2	-1.1	
		60	87.3	12.7	-1.10	2.6	61.9	-14.1	
	D8000	46	86.3	13.7	-1.00	3.8	42.3	2.0	
55		87.3	12.7	-1.00	4.1	54.5	-6.2		
60		89.2	10.8	-1.04	3.0	83.6	-25.1		
65		91.3	8.7	-1.05	3.0	114.5	-43.3		
Samsung-Plasma	D6500	51	77.7	22.3	-1.27	0.4	14.5	22.3	
		59	79.5	20.5	-1.35	0.4	28.6	7.4	
	D7000	51	81.7	18.3	-1.15	1.1	26.1	11.3	
		59	83.9	16.1	-1.22	1.0	45.1	-3.9	
		64	85.3	14.7	-1.20	1.3	55.5	-13.0	
	D8000	51	83.4	16.6	-1.06	1.4	30.2	7.6	
		59	85	15	-1.11	1.7	45.9	-3.5	
		64	87.2	12.8	-1.13	1.5	66.4	-17.2	
	LG	LW5600	47	83.3	16.7	-0.90	5.2	13.5	17.6
55			85	15	-0.96	5.0	29.0	9.2	
LW6500		47	83.9	16.1	-0.92	3.3	17.4	16.5	
		55	86.3	13.7	-1.03	2.7	46.9	0.4	
		65	89.7	10.3	-1.01	4.4	85.1	-22.0	
Panasonic		ST30	42	80	20	-0.82	3.5	2.0	26.1
	46		80.7	19.3	-0.91	3.3	6.6	23.2	
	50		80.7	19.3	-0.94	3.0	9.5	20.3	
	55		82.8	17.2	-1.04	3.0	23.0	13.5	
	60		84.2	15.8	-1.07	3.0	33.1	6.0	
	65		86.9	13.1	-1.11	2.6	59.9	-10.8	
	GT30	50	81.7	18.3	-1.12	1.9	21.4	17.8	
		55	83.4	16.6	-1.17	1.8	35.6	7.8	
		60	85.2	14.8	-1.22	1.5	59.4	-8.8	
		65	87.3	12.7	-1.20	1.5	79.0	-22.3	
		65	87.3	12.7	-1.20	1.5	79.0	-22.3	
	VT30	55	87.3	12.7	-1.12	1.8	74.0	-14.8	
		65	90.3	9.7	-1.08	2.9	115.1	-39.0	
	Sony	EX720	32	75.8	24.2	-0.91	0.3	5.8	33.1
			40	76.8	23.2	-0.97	0.4	8.0	27.7
			46	77.9	22.1	-1.11	0.3	14.0	21.0
			55	79.7	20.3	-1.18	0.3	24.5	9.6
			60	83	17	-1.29	0.2	44.3	-8.0
NX720		46	82.2	17.8	-1.03	0.9	21.1	12.5	
		55	84.1	15.9	-1.09	0.8	32.4	1.7	
		60	86.7	13.3	-1.11	0.7	60.2	-14.0	
HX820		46	85.3	14.7	-0.93	1.9	32.4	6.0	
		55	87.1	12.9	-0.95	1.9	51.0	-5.3	
HX929		46	89.1	10.9	-0.93	2.1	71.1	-13.0	
		55	90	10	-0.91	2.4	74.1	-18.0	

Table 10: Sensitivity analysis for $\rho = 2.14e - 4$. Columns (4) and (5) show the flows of consumers who in absence of a price increase would have bought the TV within one month of their arrival. Columns (6) and (7) are own-price elasticities and market shares predicted by the model with intertemporal substitution. Columns (8) and (9) show how much predicted own-price elasticities and market shares in the Buy Now counterfactual deviate from the model with intertemporal substitution.

			Coefficient	Std. Err.
Screen Size	Screen Size	mean	0.0074	0.0007
		std. dev.	0.0046	0.0002
Brand Dummies	Samsung	mean	7.9158	0.0486
		std. dev.	0.2037	0.0154
	LG	constant	8.0676	0.0401
	Panasonic	constant	7.9955	0.0405
	Sony	mean	7.6339	0.0723
std. dev.		0.3394	0.0342	
Series Dummies	Samsung D6500	constant	0.0354	0.0265
	Samsung D7000	constant	0.1491	0.0239
	Samsung D8000	constant	0.2046	0.0233
	Samsung Plasma	constant	-0.1059	0.0146
	LG 6500	constant	0.0447	0.0195
	Panasonic GT30	constant	-0.0163	0.0175
	Panasonic VT30	constant	0.0987	0.0216
	Sony NX720	constant	0.0947	0.0367
	Sony HX820	constant	0.1958	0.0358
	Sony HX929	constant	0.2935	0.0349
ϵ Std. Dev.	σ_ϵ	constant	0.1823	0.0095

Table 11: Estimates of the preference parameters and asymptotic standard errors with an annual discount factor of 0.9 without using thresholds 1 cent below the current price. The left out series dummies are the entry models of each brand: Samsung D6400, LG 5600, Panasonic ST30 and Sony EX720. The coefficient on the LG and Panasonic brand dummies are constant because the variance estimate is zero if the model is estimated with random coefficients.

Brand	Series	Screen Size	Buy the same TV after the first month (in %)	Buy a different TV (in %)	Own-price Elasticity with i.s.	Market Shares with i.s. (in %)	Buy Now: Elasticity Deviation (in %)	Buy Now: Market Share Deviation (in %)	
Samsung-LED/LCD	D6400	40	69.9	30.1	-1.96	0.7	7.5	50.7	
		46	70.2	29.8	-2.06	0.8	14.1	38.8	
		55	72.1	27.9	-2.26	0.7	28.5	11.5	
		60	74.5	25.5	-2.29	0.6	46.1	-12.6	
	D6500	40	71.5	28.5	-2.02	0.7	14.2	40.8	
		46	71.9	28.1	-1.98	1.1	17.2	33.2	
		55	73.3	26.7	-2.10	0.9	28.4	10.8	
		60	75.9	24.1	-2.19	0.7	49.3	-13.8	
	D7000	46	76.6	23.4	-1.80	2.7	27.5	14.2	
		55	78.2	21.8	-1.78	3.2	40.2	-1.7	
		60	80.4	19.6	-1.83	2.9	58.9	-21.3	
	D8000	46	78.9	21.1	-1.71	3.8	40.2	2.7	
		55	80.9	19.1	-1.66	4.7	52.0	-10.1	
		60	83.1	16.9	-1.74	3.3	79.2	-36.4	
		65	85.7	14.3	-1.77	3.3	104.8	-58.2	
Samsung-Plasma	D6500	51	67.6	32.4	-2.22	0.4	12.2	40.4	
		59	70	30	-2.32	0.4	24.6	13.2	
	D7000	51	72.3	27.7	-2.05	0.9	23.2	19.8	
		59	75.1	24.9	-2.13	0.9	39.3	-5.6	
		64	77.5	22.5	-2.06	1.3	46.3	-18.1	
	D8000	51	74.6	25.4	-1.88	1.3	26.5	13.0	
		59	77	23	-1.91	1.7	42.4	-5.3	
		64	80.1	19.9	-1.92	1.5	59.3	-25.3	
	LG	LW5600	47	75.2	24.8	-1.60	4.5	11.0	29.3
55			77	23	-1.67	4.8	25.1	14.4	
LW6500		47	76.9	23.1	-1.53	3.8	14.8	25.6	
		55	79.4	20.6	-1.71	3.0	42.8	0.0	
		65	84.4	15.6	-1.62	5.4	76.8	-30.9	
Panasonic		ST30	42	71.3	28.7	-1.49	2.9	-1.9	45.1
	46		71.8	28.2	-1.65	2.7	3.0	39.6	
	50		72	28	-1.68	2.7	5.0	34.1	
	55		73.8	26.2	-1.85	2.6	19.1	22.0	
	60		75.7	24.3	-1.85	2.8	28.6	9.7	
	65		79.2	20.8	-1.91	2.2	53.3	-15.6	
	GT30	50	72.8	27.2	-1.97	1.6	17.5	29.1	
		55	74.3	25.7	-2.03	1.5	33.8	11.9	
		60	76.6	23.4	-2.13	1.3	52.9	-13.1	
		65	79.2	20.8	-2.08	1.3	68.5	-31.7	
		65	79.2	20.8	-2.08	1.3	68.5	-31.7	
	VT30	55	79.2	20.8	-1.92	1.7	69.0	-23.0	
		65	83.6	16.4	-1.84	3.0	96.5	-51.8	
	Sony	EX720	32	68.3	31.7	-1.58	0.3	-2.8	57.8
			40	67.3	32.7	-1.72	0.4	2.1	47.3
			46	69	31	-1.90	0.4	10.6	35.3
			55	70.7	29.3	-1.98	0.4	19.1	15.8
			60	75.2	24.8	-2.11	0.3	39.7	-10.5
NX720		46	73.2	26.8	-1.80	0.8	19.3	21.7	
		55	75.8	24.2	-1.81	0.9	30.4	3.7	
		60	78.9	21.1	-1.90	0.7	51.6	-19.7	
HX820		46	77.8	22.2	-1.61	1.8	28.2	10.1	
		55	80.2	19.8	-1.61	1.9	42.7	-7.4	
HX929		46	82.8	17.2	-1.56	2.3	66.7	-20.0	
		55	85.3	14.7	-1.44	3.2	70.6	-25.6	

Table 12: Robustness check without using thresholds 1 cent below the current price. Columns (4) and (5) show the flows of consumers who in absence of a price increase would have bought the TV within one month of their arrival. Columns (6) and (7) are own-price elasticities and market shares predicted by the model with intertemporal substitution. Columns (8) and (9) show how much predicted own-price elasticities and market shares in the Buy Now counterfactual deviate from the model with intertemporal substitution.

Brand	Series	Screen Size	Buy the same TV after the first month (in %)	Buy a different TV (in %)	Own-price elasticity with i.s.	Market Shares with i.s. (in %)	Buy Now: Elasticity Deviation (in %)	Buy Now: Market Share Deviation (in %)	
Samsung-LED/LCD	D6400	40	74.5	25.5	-1.40	0.8	9.6	32.5	
		46	75	25	-1.48	0.9	14.7	25.5	
		55	77.1	22.9	-1.64	0.8	30.1	6.2	
		60	79.1	20.9	-1.69	0.7	46.4	-10.6	
	D6500	40	75.1	24.9	-1.47	0.7	15.9	27.7	
		46	76.1	23.9	-1.48	1.0	19.8	22.3	
		55	77.1	22.9	-1.58	0.7	30.0	6.1	
		60	79.6	20.4	-1.70	0.4	46.5	-12.2	
	D7000	46	80.4	19.6	-1.33	2.7	28.9	9.6	
		55	81.6	18.4	-1.36	3.0	35.8	-0.5	
		60	84	16	-1.40	2.6	58.1	-16.9	
	D8000	46	82.5	17.5	-1.27	3.8	37.0	2.4	
55		84.1	15.9	-1.29	4.2	48.1	-7.9		
60		86.2	13.8	-1.35	2.9	73.7	-28.1		
65		88.6	11.4	-1.35	3.0	102.2	-47.5		
Samsung-Plasma	D6500	51	72	28	-1.60	0.4	11.0	26.5	
		59	74.2	25.8	-1.73	0.5	23.1	9.9	
	D7000	51	77.1	22.9	-1.47	1.0	20.4	14.3	
		59	79.7	20.3	-1.54	0.9	37.7	-3.6	
		64	81.3	18.7	-1.51	1.4	46.3	-12.1	
	D8000	51	79.1	20.9	-1.37	1.9	26.6	9.0	
59		81.5	18.5	-1.43	2.0	40.5	-5.4		
64		83.7	16.3	-1.47	1.8	56.4	-20.2		
LG	LW5600	47	79.5	20.5	-1.15	5.2	12.0	20.1	
		55	81.4	18.6	-1.21	5.3	26.0	10.1	
	LW6500	47	80.1	19.9	-1.19	2.2	17.3	18.6	
		55	82.5	17.5	-1.31	1.7	41.0	1.2	
		65	87	13	-1.28	4.6	75.8	-24.8	
Panasonic	ST30	42	75.4	24.6	-1.04	3.6	-0.1	31.2	
		46	76	24	-1.15	3.4	5.8	27.5	
		50	76.5	23.5	-1.18	3.1	7.0	23.6	
		55	78.5	21.5	-1.32	3.2	20.9	15.3	
		60	80.1	19.9	-1.35	3.2	28.3	7.1	
	GT30	50	77.7	22.3	-1.42	2.0	19.5	20.0	
		55	79.4	20.6	-1.48	1.9	33.0	8.6	
		60	81.5	18.5	-1.55	1.6	52.9	-9.9	
		65	84.1	15.9	-1.56	1.5	69.8	-26.6	
	VT30	55	83.9	16.1	-1.43	1.3	67.5	-16.8	
		65	87.6	12.4	-1.39	3.1	99.2	-42.4	
	Sony	EX720	32	72.3	27.7	-1.24	0.2	0.7	41.3
			40	71.3	28.7	-1.27	0.4	6.8	34.4
			46	73.5	26.5	-1.44	0.3	13.3	24.9
			55	75.1	24.9	-1.48	0.3	23.3	10.8
60			79.2	20.8	-1.63	0.2	41.1	-9.8	
NX720		46	77.7	22.3	-1.31	0.9	20.8	14.8	
		55	80.1	19.9	-1.36	0.9	29.7	2.4	
		60	83.5	16.5	-1.46	0.7	51.4	-17.7	
HX820		46	81.5	18.5	-1.18	1.9	29.7	6.9	
		55	83.6	16.4	-1.21	2.0	41.7	-5.1	
HX929		46	86.1	13.9	-1.25	2.0	64.6	-17.9	
		55	87.6	12.4	-1.16	2.5	71.6	-22.0	

Table 13: Robustness check with nonuniform arrival. Rather than weighting all days during the sample period equally, days are weighted by the arrival rate which is estimated. The results are very close to the baseline case with uniform weighting in Table 5.

Columns (4) and (5) show the flows of consumers who in absence of a price increase would have bought the TV within one month of their arrival. Columns (6) and (7) are own-price elasticities and market shares predicted by the model with intertemporal substitution. Columns (8) and (9) show how much predicted own-price elasticities and market shares in the Buy Now counterfactual deviate from the model with intertemporal substitution. The effect on consumer surplus is unchanged compared to the baseline case with uniform weighting, i.e. an increase of \$104 or 1.7%. 39

References

- Alili, Larbi, and Andreas E Kyprianou.** 2005. “Some Remarks on First Passage of Lévy Processes, the American Put and Pasting Principles.” *The Annals of Applied Probability*, 15(3): 2062–2080.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. “Automobile Prices in Market Equilibrium.” *Econometrica: Journal of the Econometric Society*, 841–890.
- Bulow, Jeremy I.** 1982. “Durable-goods monopolists.” *The Journal of Political Economy*, 314–332.
- Carranza, Juan Esteban.** 2007. “Estimation of Demand for Differentiated Durable Goods.” *Unpublished Manuscript, University of Wisconsin-Madison*.
- Coase, Ronald H.** 1972. “Durability and Monopoly.” *Journal of Law and Economics*, 15: 143.
- Conlon, C.T.** 2010. “A Dynamic Model of Costs and Margins in the LCD TV Industry.” *Unpublished manuscript, Yale University*.
- Erdem, Tülin, Susumu Imai, and Michael P Keane.** 2003. “Brand and Quantity Choice Dynamics under Price Uncertainty.” *Quantitative Marketing and Economics*, 1(1): 5–64.
- Gowrisankaran, Gautam, and Marc Rysman.** 2012. “Dynamics of Consumer Demand for New Durable Goods.” *Journal of Political Economy*, 120(6): 1173 – 1219.
- Heckman, J.J., and B.E. Honore.** 1990. “The Empirical Content of the Roy Model.” *Econometrica: Journal of the Econometric Society*, 1121–1149.
- Hendel, I., and A. Nevo.** 2006. “Measuring the Implications of Sales and Consumer Inventory Behavior.” *Econometrica*, 74(6): 1637–1673.
- Hendel, Igal, and Aviv Nevo.** forthcoming. “Intertemporal Price Discrimination in Storable Goods Markets.” *American Economic Review*.
- Lee, Robin S.** forthcoming. “Vertical Integration and Exclusivity in Platform and Two-Sided Markets.” *American Economic Review*.
- Melnikov, Oleg.** 2013. “Demand For Differentiated Durable Products: The Case Of The U.S. Computer Printer Market.” *Economic Inquiry*, 51(2): 1277–1298.
- Mordecki, E.** 2002. “Optimal Stopping and Perpetual Options for Lévy Processes.” *Finance and Stochastics*, 6(4): 473–493.

- Nair, Harikesh.** 2007. "Intertemporal Price Discrimination with Forward-looking Consumers: Application to the US Market for Console Video-Games." *Quantitative Marketing and Economics*, 5(3): 239–292.
- Nevo, Aviv.** 2000. "Mergers with Differentiated Products: The Case of the ready-to-eat Cereal Industry." *The RAND Journal of Economics*, 395–421.
- Nosal, Kathleen.** 2012. "Estimating Switching Costs for Medicare Advantage Plans." *Unpublished manuscript, University of Mannheim.*
- Petrin, Amil.** 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*, 110(4): 705–729.
- Rust, John.** 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica: Journal of the Econometric Society*, 999–1033.
- Schiraldi, Pasquale.** 2011. "Automobile Replacement: A Dynamic Structural Approach." *The RAND Journal of Economics*, 42(2): 266–291.
- Shcherbakov, Oleksandr.** 2013. "Measuring Consumer Switching Costs in the Television Industry." *Unpublished manuscript, University of Mannheim.*
- Song, Inseong, and Pradeep K Chintagunta.** 2003. "A Micromodel of New Product Adoption with Heterogeneous and Forward-looking Consumers: Application to the Digital Camera Category." *Quantitative Marketing and Economics*, 1(4): 371–407.
- Stokey, N.L.** 1979. "Intertemporal Price Discrimination." *The Quarterly Journal of Economics*, 355–371.
- Su, Che-Lin, and Kenneth L Judd.** 2012. "Constrained Optimization Approaches to Estimation of Structural Models." *Econometrica*, 80(5): 2213–2230.