Learning by Doing and Consumer Switching Costs*

Yufeng Huang†

November 1, 2015

Abstract

This paper shows that using a product builds up specific human capital that is difficult to transfer to another product, and this creates a form of consumer switching cost. In the context of digital cameras, I use novel data to directly characterize the cost of switching products, by measuring the changes in consumer picture quality. I estimate a structural model of demand with learning by doing, in order to quantify the effect of learning-driven switching cost on the demand for new products and on consumer welfare. I find that 10% of consumer human capital is lost to every brand switching, for which a consumer is willing to pay $40 to avoid. This explains a quarter of the consumer's persistence in brand choice, and lowers her cross price elasticities by a factor of 3.

Keywords: human capital, learning by doing, switching cost, dynamic programming

*This paper comes from the first chapter of my PhD dissertation in Tilburg University. I am indebted to my advisors, Bart Bronnenberg and Tobias Klein, and also to my committee member Jaap Abbring. I thank Guy Arie, Pradeep Chintagunta, Jonne Guyt, Günter Hitsch, Jiaxuan Li, Mitch Lovett, Chris Nosko, Peter Rossi, Raphael Thomadsen and Yifan Yu, and many seminar and conference participants for their feedback. I thank Marketing Science Institute (MSI) for their acknowledgment in the 2013 Clayton Competition. I thank The Netherlands Organization for Scientific Research (NWO Vici Grant 453-09-004) for financial support. All errors in this paper are mine.

†Assistant Professor of Marketing, University of Rochester Simon Business School. Email: yufeng.huang@simon.rochester.edu.

1 Introduction

> "I've been using Dvorak (keyboard) since 5.5 years now, [...] If you decide to switch, do it the hard way. Switching is terrible the first days."
>
> – replies to *Thinking of Switching to Programmer Dvorak Layout*[1]

> "Nikon and Canon are as good as each other overall. [...] The differences lie in ergonomics and how well each camera handles, which is what allows you to get your photo – or miss it forever. [...] and I can't for the life of me figure out the menus of the Nikon Coolpix cameras."
>
> – Ken Rockwell, *Nikon vs. Canon*[2]

A consumer who has learned how to use one product is unwilling to switch to another, if it handles and operates very differently from what she has learned. This is because her human capital on the previous product does not apply to the new one, therefore she has to learn again. For example, David (1985) documents the path dependence of the modern keyboard layout – "QWERTY" – despite alternatives such as the Dvorak layout have been shown to be superior. This and many other examples show that product-specific learning by doing gives rise to a form of consumer switching cost.

In this paper, I directly measure the extent to which consumer human capital – developed through repeatedly using one product – is non-transferable to other products. Such learning cost is a predominant example of consumer switching cost (David, 1985; Klemperer 1987, 1995), but has not been measured empirically. This is because, existing measures of switching costs rely on the persistence in consumer choices (Shcherbakov, 2008; Dubé et al., 2010), which could be confounded with unobserved tastes (Heckman, 1981), and does not distinguish between different sources of switching costs (Greenstein, 1993). In the context of digital cameras, I address this concern utilizing novel data on pictures produced by a set of camera users. For these users, I observe

[1] Extracted from the following URL in October 2015. http://stackoverflow.com/questions/4264853/thinking-of-switching-to-programmer-dvorak-layout

[2] Extracted from the following URL in March 2014. http://www.kenrockwell.com/tech/nikon-vs-canon.htm

their camera purchase decisions, and construct a metric for their picture quality. This allows me to directly characterize learning by doing as well as consumer switching costs, respectively, by the observed changes in picture quality as a result of using a camera or switching to a new camera. I then estimate a dynamic structural model of durable goods demand with learning by doing, in order to quantify the role of learning-driven switching cost on the choice for new products, on price elasticities, and on consumer welfare.

My information on picture quality and camera usage history comes from picture level data, collected from Flickr.com. Camera and picture-taking time information is automatically recorded by the digital camera, and displayed on the website. This gives long individual-level product usage histories, with a median observation window of 6 years. However, a viewer on the website does not observe this information until she clicks to view an enlarged version of the picture. This means, if pictures taken by different cameras or at different points in time are uploaded in close sequence, differences in their cumulative number of views are indicative of the underlying picture quality. I take this revealed viewer-preference approach to construct an implied picture quality metric, essentially as *unexplained* differences in views among pictures that are uploaded together, and share the same topic. I also perform robustness checks using alternative measures of picture quality.

Descriptive patterns in the data strongly suggest that a considerable part of consumer knowledge is product specific. While picture quality is higher for consumers who have more experience in general, the quality of pictures produced by their recently-adopted cameras are lower, *proportional* to their experience. Also, picture quality sharply declines following camera replacements, but is recovered in a few months. This suggests that while the ability to take good pictures can be trained over time, a part of it is tied to using a specific product. This part of product-specific knowledge pushes up consumer switching costs over time. In addition, I find that across-brand switching decisions are less common for an experienced user, compared to her camera replacement choices within a brand. The descriptive analysis provides qualitative evidence that, for knowledge-intensive durable goods, product specific experience take up a large share of a consumer's knowledge stock, and shape her choices of new products towards familiar brands.

I construct a structural model to quantify the impact of product-specific learning by doing on consumer demand for new products. The model brings classical learning by doing and human capital formation theory (Becker, 1965; Michael, 1973) to the context of dynamic demand for durable goods (Melnikov, 2013; Song and Chintagunta, 2003; Gowrisankaran and Rysman, 2012), and allows one to separate product-specific learning by doing from alternative explanations of switching costs. In the model, a consumer uses her camera to produce picture quality through a production function (Michael, 1973), in which her human capital complements the camera quality. Using the camera also allows her to learn by experimentation, but a part of the added human capital is tied to the type of camera she uses, as well as the way it is designed by the manufacturer. Hence, when switching to a different camera, the product-specific part of her human capital is lost – in addition to other switching costs she might have. She makes forward-looking camera replacement decisions, rationally expecting the potential consequences of her current choices. To make the model realistic, I also allow for differences across consumers, in their (time-invariant) preferences, their first cameras and initial knowledge, as well as the way past history affects their current decisions. Finally, with data on pictures created by different generations of cameras, I can capture technology evolution in a flexible yet simple way.

The primary contribution of this paper is that I show how one can directly quantify the magnitude of consumer switching cost. I measure changes in picture quality in connection to a structural model of camera demand, and find that switching between brands costs 10% of a consumer's human capital, for which she is willing to pay $40 to avoid. This learning cost is 30% of the average price of a compact camera, or 7% of that of a DSLR; this is comparable to the findings in Dubé et al. (2010) for consumer packaged goods, but is smaller than in Shcherbakov (2008) for cable television contracts. Different from these papers, the learning cost I quantify comes from a specific channel, and is not evaluated using persistence in choices. Therefore, it should be understood as a fraction of the total consumer switching costs.

Relatedly, I find that brand-specific learning by doing alone can explain 22% of the path dependence of an experienced consumer's brand choice. Recent empirical literature document that shocks unrelated to a consumer's taste can have persistent effects on her brand choices (Bronnenberg et al., 2012), and this paper provides a specific explanation to why this is the case. This implies that brand-specific learning cost is an important explanation to advantages for the first-mover (David, 1985; Bronnenberg et al., 2009). I also find that cross price elasticities would have been 3 times as large if there were no brand-specific learning cost. This suggests high ex-post market power to a brand's own customer base (Klemperer, 1987; Gabszewicz et al., 1992).

Finally, this paper provides new evidence to consumer learning. In my context, I document that consumers become better at using a product because they are able to reproduce good methods, discovered from past experience. This complements the existing literature on consumer learning, which primarily emphasize that consumers learn by reducing uncertainty (Erdem and Keane, 1996; Crawford and Shum, 2005) or by correcting wrong prior beliefs (Israel, 2005).

The remainder of the paper is structured as follows. Section 2 gives a brief review to the literature related to this study. Section 3 describes the data collection process and how I define the key variables – in particular, the identification strategy that allows us to measure picture quality. Section 4 then presents descriptive evidence for consumers' general and product-specific learning by doing. Given the evidence, Section 5 outlines an empirical model of experience evolution and consumer choices on purchasing and using cameras. Next, Section 6 presents and discusses model estimates, elasticities and fit. Section 7 then quantifies the role of brand-specific learning by doing, through counterfactual experiments. Finally, Section 8 concludes.

2 Related literature

This paper is positioned in the intersection of the consumer human capital literature and the switching costs literature. Among the switching cost theory literature, Klemperer (1987; 1995) discusses various sources of consumer switching costs, and shows that firms balance between setting high prices to exploit locked-in consumers, and setting low prices to get new consumers.[3] The consumer human capital literature is small and has been mostly theoretical: Michael (1973) discusses the implication of consumer human capital – in a framework related to Becker (1965) – but does not find convincing evidence using education as a proxy. Ratchford (2001) discusses the implication of this consumer human capital on firm behavior and optimal managerial decisions. Jovanovic and Nyarko (1996) presents a Bayesian version of learning by doing, and characterizes a consumer's technology adoption decisions when upgrading to more advanced technology incurs switching costs.

[3]Farrell and Shapiro (1988), Arie and Grieco (2014) amd Dubé et al. (2009) show conditions under which prices are lower in presence of switching costs.

The primary contribution of this paper is that I propose an alternative way to measure switching costs. Conventionally, switching cost is empirically measured using un-explained persistence in consumer choice. Bronnenberg et al. (2012) document that consumer's brand preference is persistent geographically, shown up in the slow adjustment of brand shares for consumers who migrate across states. They are not specific in pinning down what type of switching cost it is, and the associated implications. Shcherbakov (2008) estimates the magnitude of transaction cost in the cable TV industry. Dubé et al. (2010) estimates the inertia in consumer choices, controlling for flexible specification of heterogeneity; they point out that the inertia can be viewed as switching costs. Crawford and Shum (2005) document the presence of uncertainty for physicians in the prescription drug market, reflected in their unwillingness to switch to new drugs. Nosal (2012) finds large costs in switching between medical insurance providers, utilizing exogenous shocks that pushes a household in search of new providers. In this paper, I directly measure the impact of switching cost – specifically, the consumer learning cost – by changes in the quality of pictures for camera users. In this way, one is able quantify switching cost without confounding it with unobserved tastes. Also, this direct measure isolates product-specific learning by doing from alternative explanations of switching costs,[4] thus has direct policy implications.

This paper is also related to the literature on consumer learning cost. David (1985) provides a historical account of why the QWERTY keyboard survives and dominates the market. Greenstein (1997) studies the case of government procurement of mainframe computer in the 70s and 80s, and documents that buyers anticipate, and take actions to prevent, potential incompatibility issues in the computer systems they adopt. Greenstein (1993) estimates a simple logit model about computer buyer's choice of supplying firm, characterizing the observed persistence in the choice of supplier as a function of the duration of existing partnerships. His reasoning is that learning cost increases with the partnership duration, while any additional persistence in the choice of partner could be other types of switching cost. My paper follows similar logic to Greenstein (1993) but provides more direct evidence of consumer learning costs. Gabszewicz et al. (1992) present a theoretical framework where consumers are heterogeneous in their cost of learning to use a product. They emphasize that it is important to distinguish whether such learning cost is brand-specific, since

[4]For example, relying on partnership duration variation could confound with learning and uncertainty (Crawford and Shum, 2005).

general knowledge on one firm makes it easier to switch to another firm for the consumer, while brand-specific knowledge makes it harder. They also characterize prices and entry decisions under brand-specific learning, and show that brand-specific learning might increase prices but not entry barrier, and hence has different implications to other types of switching costs.

In terms of methodology, the consumer demand framework in this paper is derived from the literature on dynamic discrete choice of differentiated products. Using aggregate data, Melnikov (2013) estimates the demand for durable goods, of forward-looking consumers who are only in the market once. Song and Chintagunta (2003) applies this framework to the digital camera market. Gowrisankaran and Rysman (2012) extends this framework to allow for replacement decisions, but impose a dimensionality-reduction assumption on the state space, to ease the computational burden. On the other hand, with access to individual level decision data, Hendel and Nevo (2006), among others, extend the dynamic discrete choice framework of Rust (1987) and estimate demand for differentiated products. My model treats differentiation in product quality in a different way, compared to Hendel and Nevo (2006) and Gowrisankaran and Rysman (2012). In addition, my paper jointly models usage and purchase decisions, similar to Albuquerque and Nevskaya (2012).

Finally, it is important to contrast this paper with the literature on demand estimation with consumer learning. Erdem and Keane (1996) estimates the demand of forward-looking consumers, when they face uncertainty of product quality, which is gradually resolved by repeated purchase. Crawford and Shum (2005) model physician's attitude towards new drugs in a similar way. Osborne (2007) exploits the different locked-in patterns, and estimates a demand system with brand-specific learning and instantaneous switching costs.[5] Lovett and Staelin (2012) model TV viewing when viewers learn about their own taste, and separately identify belief and utility through differences in stated and revealed preference. Li (2015) estimates individual demand of digital cameras, when consumers learn about their general taste (towards photography in this case), but face switching costs between brands. My paper provides evidence that consumers learn because they are able to discovered and retain good methods of using a product, thus accumulate product-specific human capital by using. This is a different mechanism compared to resolving consumer's uncertainty through Bayesian updating.

[5]In a sense, his idea is related to Greenstein (1993).

3 Data

3.1 Collection

I extract picture level data from Flickr.com – a popular photo sharing website. The data extraction was implemented between March 2012 and April 2013, until a major change in user-interface took place on Flickr. For each individual account, the picture data dates back to when the account was established on Flickr – as early as 2000. During the data collection period, pictures (including their detailed information) were visible for all viewers, with or without an account.

Camera-recorded information is embedded in each picture, as *Exif* (exchangeable image file format) data. For the purpose of this paper, those data contain valuable information for camera identity, as well as the date when the picture was taken. To complement the Exif data, I also collect information on the date of upload, and the cumulative number of views and "favorite" votes, from the time of upload, until data extraction. Figure 1 provides an example of the information I extract for one picture.

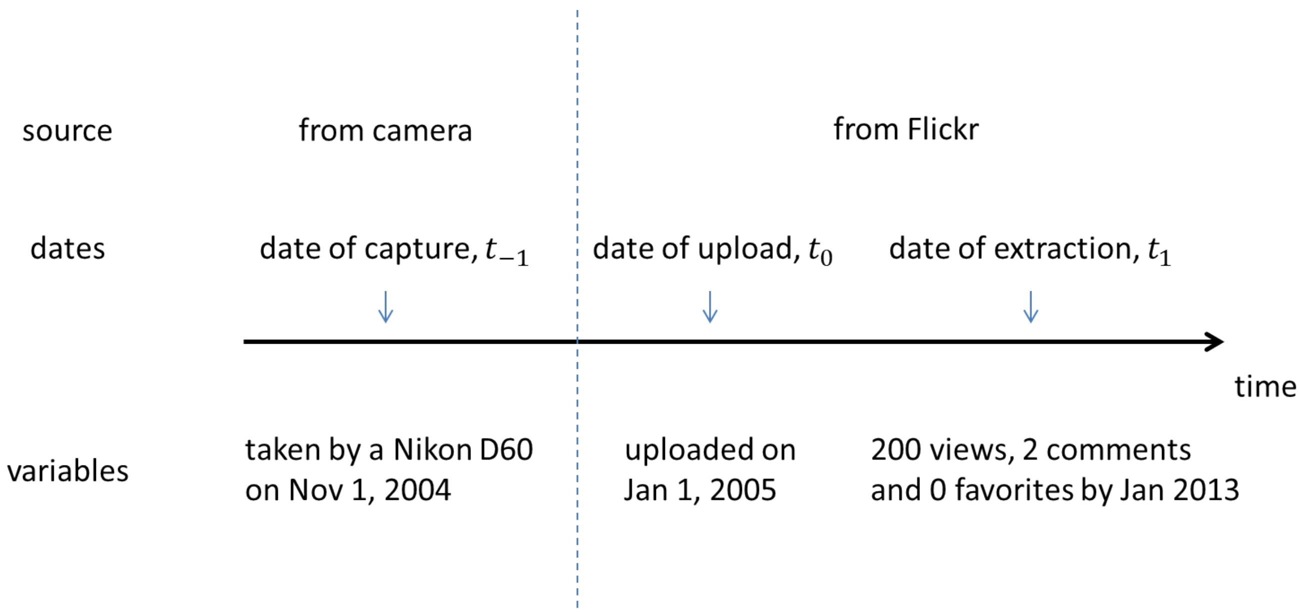

Figure 1: Structure of picture-level data from Flickr: an example

Note: This figure shows the structure of the picture level data, extracted from Flickr.com. From the camera-recorded data (Exif), I collect the camera identity (in this example, Nikon D60) and capture date. From Flickr-recorded data, I collect the upload date, as well as the cumulative views and "favorite" votes from upload to data-extraction.

I collect data on two levels. On the picture level, I sort an individual's pictures in the order

Table 1: Sample Selection Criteria

	percentage
Taken by compact camera or DSLR	96.4
Exif data complete	89.6
Taken after year 2000	89.0
All above criteria	75.6
obs.	2777728

Notes: This table reports the sample selection criteria. On the picture level data, I drop observations on pictures taken by camera formats other than compact camera or DSLR, or with Exif data that lack an indicator of camera model or date of capture, or taken before January 1, 2000. Altogether, this excludes 24.4% of the sample.

of upload dates, and collect the data *once*, from one in every five pictures. This gives me *cross-sectional* data on picture level information.[6] On the individual level, I collect data on Flickr-summarized monthly picture-taking and uploading records, for each individual.

I also gather cross-sectional data on camera characteristics, and longitudinal Ebay auction price data. The camera characteristics data is compiled from the Flickr camera database, DPreview.com, and Cnet.com. In addition, Pixel-peeper.com summarizes monthly history of Ebay average auction prices per camera, from late 2006 onward.

3.2 Sample selection and summary statistics

I focus on the users with paid accounts – called the "pro accounts" – at the time of data collection.[7] These accounts allow users to upload many more pictures than the free accounts, creating more precise measure of both their camera-switching decisions, and their human capital growth path (explained later).

I collect data from all paid account users with a username no longer than 5 letters/digits. Focusing on shorter usernames gives me users with long histories on Flickr,[8], assuming that usernames are exogenous to a user's unobserved preferences.

Sampling one in every 5 pictures gives me close to 2.8 million observations on the picture level. Among these data, I discard the pictures taken by cell phones, film cameras, camcorders or

[6] I did not re-visit a picture multiple times, because the time spent on collecting data from each picture is large.

[7] Flickr offers either a free account – which is imposed a monthly upload capacity as well as a maximum-viewable-pictures restriction (200 pictures in total), or a "pro account" that costs $24.95 (as in 2012) annually.

[8] The in-sample duration is not orthogonal to preferences and choices, and hence I do not condition on users with long in-sample duration.

Table 2: User Level Data Summary

	Mean	Median	Stdev
months since registered in Flickr	69	74	24
number of contacts at data extraction	94	20	292
total number of pictures	1691	981	1897
number of in-sample pictures	359	203	410
number of cameras ever used in-sample	4	3	4
max views per month, first 10 pic	7	1	58
max views per month, last 10 pic	20	4	193
price of the least expensive camera used	216	157	191
price of the most expensive camera used	1040	762	655
obs.	5499	5499	5499

Notes: The table reports summary statistics from the data. Mean, Median and StDev are the mean, median and standard deviation of the data, respectively. The number of contact is the number of other accounts, who are followed (subscribed) by the given user at the time of data extraction. The number of cameras ever used in-sample is the number of unique camera identities one observes from the user's Exif data. Prices of the least and most expensive cameras are in 2005 US dollars.

digital media players, or those claimed to be taken prior to year 2000 (which is more likely to be a mistake in the camera date settings), or have incomplete Exif data (in particular when identities of the cameras or the picture taking time are missing). This excludes 24.4% of the picture data – as shown in Table 1.

Table 2 provides summary statistics for the user level data after sample selection. Note that the median duration of observation for a user is beyond 6 years, giving a long observation window for the slow evolution of photography knowledge. Also, there is a considerable increase in the maximum views per unit time among pictures taken at the beginning of the sample, compared to those taken at the end of the sample; while the views have a larger spread towards the end of the sample. This suggests both an *increase* and a *divergence* in the number of views one's pictures can attract.[9] Besides, in the median, an individual only subscribed to 20 other users. Compared to Facebook users, this suggests that the increase of views might not be driven by social-network changes.[10] Finally, in the median, an individual has used 3 cameras, with considerable dispersion in the prices at the time of first-usage: the real (Ebay auction) price of her most expensive camera is more than twice of the price of her least expensive camera.

[9]Which might be due to changes in picture quality, or changes in the size of user base of Flickr.com.

[10]As a comparison, the median Facebook user has 200 friends, by account of Aaron Smith (extracted in June 2014, from http://www.pewresearch.org/fact-tank/2014/02/03/6-new-facts-about-facebook/).

3.3 Implied picture quality

3.3.1 Intuition

By the Flickr website setting, a viewer does not see the capture date of a picture, or the camera behind it, before she clicks on the picture link for a more detailed view. Therefore, variations in the picture-taking date and camera identity, on the total number of clicks on a picture, is informative of the picture quality. I find that controlling for the *time window* when pictures are displayed on the website, pictures taken later, as well as taken by more advanced cameras, systematically enjoy more views. Appendix A documents one version of such findings. This suggests that camera equipment and user experience – not directly observed on the web page – are "observed" through the appearance of the picture. More experienced users and users with more advanced equipment create pictures that the audience like better.

With this intuition, I construct a measure of "picture quality", defined as the residual difference in log views, between pictures uploaded in the same time window. This implied picture quality measures how much a "representative" viewer likes one picture versus another, in terms of her likelihood of clicking on either picture. In addition to controlling for the display time window, we also control for measures of different photographer's popularity, and popularity differences between the pictures' topic.

In this section, I first layout a model that generates a linear relationship between the implied picture quality, and other explanatory variables to differences in the number of views (clicks). I then discuss summary statistics of the measure. This implied picture quality will be later treated as data in the descriptive analysis and structural estimation.

3.3.2 A model of viewer clicks

I denote the cumulative number of views of picture p captured by individual i, as the accumulation of an underlying viewer-flow process to the photographer. We model the flow process $flow_{ipt}$, as a function of the underlying picture quality q_{ip}, the overall flow of viewers on Flickr.com ϕ_t, and other observed characteristics of the picture that are not related to quality, z_{ip}:

$$views_{ip} = \sum_{t_0 \leq t \leq t_1} flow_{ipt} \tag{1}$$

where

$$flow_{ipt} = \phi_t \exp\left(q_{ip} + z_{ip}\psi\right),$$

Omitting i and p subscripts, I denote t_0 and t_1 to be the calendar dates of upload and data extraction, respectively.[11] The cumulative number of views is the summation of the viewer flow between these two dates. In the viewer flow specification, q_{ip} is the (unobserved) quality of the picture, which is implicitly a function of user experience, camera, and an econometric error.[12]

Take the log of Equation (1), we have

$$\log\left(views_{ip}\right) = \Phi_{t_0 t_1} + z_{ip}\psi + q_{ip}, \tag{2}$$

where $\Phi_{t_0 t_1} = \log\left(\sum_{t_0 \le t \le t_1} \phi_t\right)$ is a time-window-specific fixed effect, that captures the overall cumulative viewer arrival in the time window $[t_0, t_1]$, when the picture was on display. Also, z_{ip} includes the topic of the picture, the number of pictures uploaded, the order of upload, and the duration since the photographer was registered on Flickr – these help control for additional variations in views that are not related to picture quality.

The specification (2) makes two assumptions. First, the upload *timing* decision of picture p is orthogonal to the unobserved quality q_{ip}, up to an individual fixed effect. That is, given an individual's fixed characteristics, she does not time the upload in the order of their quality.[13] As support of this, we find that more than 3/4 of all pictures are uploaded in the immediate next batch, which implies that the (infrequent) upload decision might be driven by other time costs. In addition, as shown in Figure 13 in the Appendix, the views on the late-uploaded pictures are not systematically different from others. This suggests that their quality is not selectively different.

As the second assumption, I impose that the cumulative viewer base, $\Phi_{t_0 t_1} = \log\left(\sum_{t_0 \le t \le t_1} \phi_t\right)$, is the same for all individuals, up to heterogeneity in individual fixed effects.[14] We relax this assumption in our reduced form tests of the return to experience – one version presented in Appendix A – and find the results to be similar.

[11] Note that t_0 and t_1 are picture specific.

[12] One might alternatively interpret this as a noisy measure of picture quality.

[13] For example, this assumption will be violated if an individual decides to upload good pictures first, and the lower quality pictures in later batches. If this is the case, the upload time t_0, and hence $\Phi_{t_0 t_1}$, will be correlated with q_{ip}.

[14] This assumption would be violated if for some individuals, their viewer base *increase* faster than others. If this is the case, we will over-state the trend in q_{ip}, and attribute it to learning effects.

I code all control variables in z_{ip} in indicator to allow for flexibility, and then estimate Equation (2) by ordinary least squares, with individual fixed effects in q_{ip}. Then, I collected the projected fixed effects and residuals. This gives a measure of quality for each individual picture p, taken by photographer i.[15]

Finally, there is a concern that the observed picture quality distribution is selected, and the selection criteria changes over time. To this end, I take the upper bound of the quality measure, among pictures taken in the same month:

$$Q_{ijt} = \max_{p \in t} \hat{q}_{ip},$$

and term it the *implied picture quality*. This alleviates the selection problem if we assume that the best possible picture is always included in the upload set. Robustness checks both in the descriptive evidence and structural model show that the moment (or order statistics) in the implied quality distribution does not change the main findings of this paper.

3.3.3 Summary statistics of the implied picture quality

Table 3 summarizes the implied picture quality in each month. First, we find that quality increases with experience but with decreasing marginal return. This is consistent with existing learning by doing evidence in the contexts of labor (Shaw and Lazear, 2008; Levitt et al., 2013) and firm (Benkard, 2000; Kellogg, 2011). Second, using a small sub-sample with non-zero favorite votes data,[16] I find that the correlation between implied picture quality and monthly maximum rating (if nonzero) is around 0.55, which cross-validates that the implied picture quality measure is well correlated with measures of likings on the picture.[17]

[15] The projected individual fixed effects are considered to be in q_{ip}, as proxy of start-of-the-sample picture quality. Because we can trace every user to her starting point in Flickr, but not to her initial experience in photography, it is unlikely that the initial conditions – reflected in individual fixed effects – are due to heterogeneity in the viewer base.

[16] The share of individual-monthly observations where *at least one* picture has received *at least one* favorite vote is 15%.

[17] We adjusted for the number of pictures uploaded in a month and find that the upload volume does not explain such correlation.

Table 3: Summary of the monthly maximum implied picture quality

	max quality	stdev	max favs	corr. with qual.	corr. adj. for nr pic
0 year of expr	0.593	1.479	0.636	0.578	0.557
1 year	0.984	1.494	0.749	0.581	0.523
2 years	1.198	1.558	0.819	0.577	0.509
3 years	1.268	1.574	0.867	0.557	0.499
4 years	1.276	1.553	0.878	0.529	0.482
5 years	1.236	1.546	0.882	0.515	0.478

Notes: This table summarizes the individual-monthly maximum of the inferred picture quality (Section 3.3), which is treated as data in the subsequent analysis. Monthly maximum refers to quality of the best picture *captured* in the given month, by an individual. Years of experience is defined as number of years from the first in-sample picture to the current month of picture-taking. The first two columns summarize its mean and standard deviation. The third column presents average of the highest rating ("favorites") one gets for pictures taken in the month, given that the highest rating is non-zero (15% of the individual-month data). The fourth column presents its correlation coefficient with the highest inferred quality. Finally, the fifth column presents this correlation adjusted for the number of pictures taken (or has favorites) in the given month.

3.4 Camera ownership

As previously mentioned, the camera identity is embedded in the picture's Exif data, which I use to infer camera ownership history. To do so, for any given individual, when a new camera appears to have taken at least two pictures, I assume that this camera is purchased in the month when the first picture was taken. I also assume that it replaces the previous camera.

For 75% of all individual-camera combinations, I never observe an old camera taking pictures after arrival of the new one. For the remaining 25%, although the earlier cameras still take at least one picture at some point in time, the majority of the pictures are taken by the most recently acquired camera. Overall, for more than 90% of the time, after a new camera – as defined – comes into the sample, we *never* observe the previous camera taking most pictures in *any* month. This is also graphically illustrated in Figure 14

3.5 Price indices

I use the monthly average Ebay auction price for each camera model, as proxy for the price one expects to pay. This price data is averaged across first and secondary markets, therefore better represents the average prices consumers face – especially for older models.

I deflate prices to 2005 US dollars. Then, separately for both camera formats, I take the average of the prices of all available cameras in a given month, weighted by their market shares in the Ebay auction data.[18] Since the data only range from 2006 onward, I interpolate the missing values before 2006, by taking a log-linear fit against time, plus a *simulated* regression error.[19]

4 Descriptive analysis

4.1 Overview

This section presents descriptive analysis on the role of consumer experience in photography. I first show that longer experience in photography is associated with higher quality pictures, reflected in the popularity of those pictures among then Flickr audiences. This demonstrates that consumers learn photography after they start taking pictures.

I then present two pieces of evidence for the learning by doing process. First, I show that part of the consumer human capital is lost when switching between cameras, which creates an increasing switching cost in the consumer experience. Second, in Section 4.4, I show that improvement in human capital is mainly achieved through learning by experimentation. Our structural model about learning by doing captures this.

Note that consumers with low human capital or low preference towards photography might systematically drop out earlier, which causes a selection problem. To address this, I first restrict attention to a sub-sample of individuals, who stays in the sample for at least 5 years.

4.2 Evolution of general and product-specific human capital

As elaborated in Section 3, I use the number of views as a proxy for picture quality. To be able to do so, I correct for confounding explanations that also changes the cumulative number of views – such as the duration of display, popularity in a topic, or growth of a user's network.

[18] That is, the number of auctions for a given camera model, as a percentage of the total number of auctions in the sample.

[19] Separately for each format, I regress log price index on a linear time trend, and interpolate the missing value using the linear prediction plus a simulated prediction error. The R-squared for the linear regression are around 0.7 for both formats. Keane and Wolpin (1994) use this method to interpolate missing data in their value function calculations.

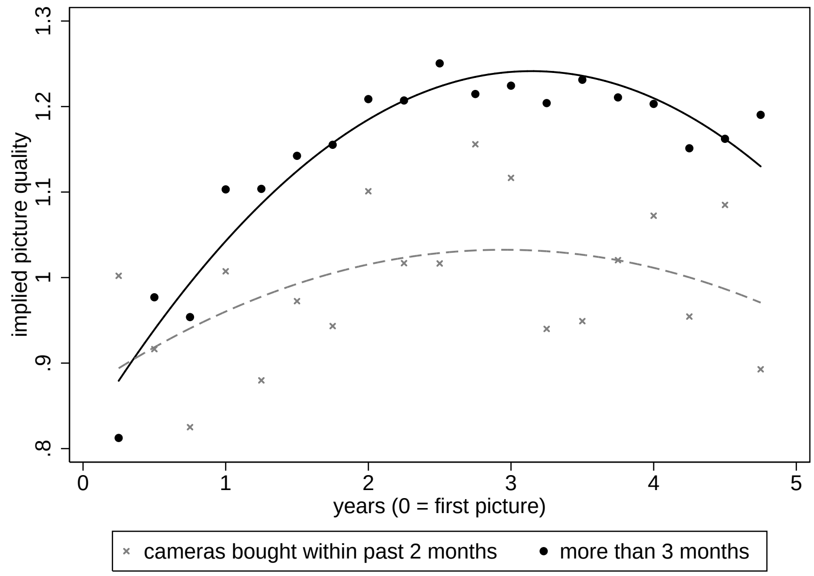

Figure 2: Evolution in picture quality, new and familiar cameras

Notes: This figure plots the implied picture quality (popularity-adjusted views) against years since the first picture on Flickr. We also separate pictures captured by cameras that the individual has been using for more than 3 months (solid circles and line), and pictures captured by cameras bought within the past 2 months (crosses and dashed line).

Figure 2 presents the evolution in the popularity-adjusted views – as a proxy for picture quality – as the experience in photography of a user grows. The increasing time trend shows that experience in photography is associated with higher picture quality, which is evidence of human capital evolution. In addition, we separately plot pictures captured by cameras that a consumer is familiar with (illustrated by solid points), and those by cameras that are new to the consumer (grey crosses). We find that as a consumer gets more experienced, the gap between pictures taken by new and familiar cameras becomes larger. This shows that part of her human capital is product-specific and cannot be transferred to the next camera.

4.3 Evolution of human capital around camera-switching

Figure 3 presents changes in picture quality around the instance of camera switching – normalized to year zero. We find that there is an immediate drop in picture quality at switching, which is recovered in 2-3 months. In addition, picture quality further gradually increases to a higher level, in the next 1.5 years. This is in line with the gap between new and familiar cameras, presented in

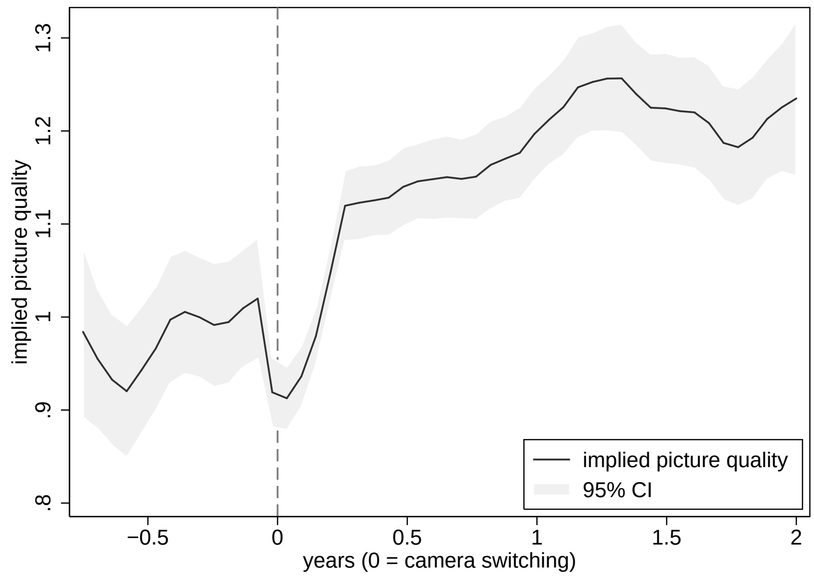

Figure 3: Quality of best pictures around camera switching

Notes: This figure depicts the changes in maximum picture quality, around the period when an individual switches cameras. We focus on the years before and after a consumer switches her camera at year 0. With the left vertical axis, the dark line (and shaded areas as its 95% confidence interval) depicts the maximum picture quality that the consumer can produce, using her old camera until year -1/12, and new camera *from year 0*. The line is estimates of local polynomial regression with bandwidth 1.

Figure 2. This suggests that at the instance of camera switching, the individual loses both explicit knowledge on camera operations (e.g. menu and button layout), as well as implicit knowledge on camera usage (e.g. how to best circumvent a certain product limitation). While the first can be quickly learned in a month or two, the second can only be learned with long experience with the new camera.

In the appendix, we show that the drop in picture quality is not explained by changes in the number of pictures – as the number of pictures goes up when one starts using a new camera. We also present robustness checks to various alternative proxies of picture quality.

4.4 Stochastic growth in consumer human capital

I next explore how consumer human capital improves. In particular, I test between whether human capital is improved through time or through picture-taking practices. The latter is an outcome of learning by doing, while the former would suggest learning happens externally, plausibly from others' knowledge.

To take a closer look at the human capital evolution, I formulate picture quality in the following specification:

$$Q_{ijm} = \theta_0 + \theta_m \cdot m + \theta_q \cdot Q_{ij'm-1} + \theta_t \cdot t_m + \theta_k \cdot SLR_{im} + \theta_i + \vartheta_{im}$$

where m is the number of months the individual has been taking pictures, t_m is the calendar month of picture-taking month m, and $Q_{ij'm-1}$ is picture quality from the previous picture-taking month $m-1$. I also control for changes in equipment SLR_{im} and individual fixed effect θ_i.[20]

One should find that θ_t explains the majority of evolution in picture quality (other than the camera coefficient θ_k), if the source of knowledge comes externally. Otherwise, either θ_m or θ_q (or both) explain changes in Q_{ijm}. In addition, steady contribution of each picture-taking month m supports hypothesis of deterministic growth of human capital, where learning comes from step-wise accumulation of knowledge through practice. On the contrary, if a strong θ_q supports the hypothesis a stochastic growth of human capital, where learning instead comes from being able to replicate shocks in past picture quality, plausibly out of experimentation.

We estimate a first-differenced version of the above specification:

$$\Delta Q_{ijm} = \theta_m + \theta_q \cdot \Delta Q_{ij'm-1} + \theta_t \cdot \Delta t_m + \theta_k \cdot \Delta SLR_{im} + \Delta\vartheta_{im}, \tag{3}$$

where we denote Δ as the first difference operator (between m and $m-1$). We instrument $\Delta Q_{ij'm-1}$ by $Q_{ij'm-2}$, because the term is correlated with first-differenced past error term $\Delta\vartheta_{im-1}$.[21] The results are presented in Table 4.

I find a positive carry-over from a change of picture quality in the most recent past, to the change of picture quality in the current month. For a 1-standard deviation shock in picture quality in $m-1$, current picture quality increases by 0.06. This is about 1/3 of the average annual growth rate of picture quality, observed in Figure 2. On the other hand, the time trend θ_t and the

[20]Implicitly, we assume that the error term is serially uncorrelated, and orthogonal to all right-hand side variables.

[21]This is because they have common component ϑ_{im-1}. We can instrument this by $Q_{ij'm-2}$ due to the assumption that ϑ_{im} is serially uncorrelated. See Arellano and Bond (1991) for details.

Table 4: Reduced-form picture quality evolution

	D.quality	D.quality, t-1 (1st stg.)
D.quality, t-1	0.076***	
	(0.006)	
D.dslr	0.187***	-0.076***
	(0.017)	(0.013)
D.time	0.002	-0.024***
	(0.002)	(0.001)
quality, t-2		-0.346***
		(0.006)
constant	0.007*	0.484***
	(0.003)	(0.011)

Note: This table presents Arellano-Bond estimates for Equation (3). "D" denotes Δ, or the first difference operator. The constant term represents the experience effect before first-differencing. The endogenous $\Delta Q_{ij'm-1}$ is instrumented by $Q_{ij'm-2}$, and the second column presents first stage estimates. Standard errors are heteroskedasticity robust and clustered by individual. Standard deviation of residuals $\Delta\vartheta_{im}$ is 1.17. This implies that the standard deviation for ϑ_{im} is 0.83.

deterministic learning by doing coefficient, θ_m, are both small and insignificant from zero at 95% confidence. Hence, we find that improvement in picture quality comes from shocks of past quality.

This finding supports the hypothesis of consumer knowledge improves via practice, and in particular, via their ability to replicate shocks in picture quality. Arguably, these shocks might come from experimentation. In line with these findings, our structural model characterizes learning by doing as the ability to preserve only the knowledge that improves current and future picture quality.[22]

4.5 Increasing tendency to switch within a brand

Finally, I document that consumers respond to changes in their human capital – general and product specific – in their choice of next camera. In Figure 4, I find that consumers switch between cameras more often, when they have accumulated more experience. Plausibly, having more experience allows them to value picture-taking activities higher, thus increasing their valuations towards the latest camera technology or more advanced cameras. In the appendix, I also document that having

[22]The symmetric structure here also implies quality-destroying knowledge further destroys future human capital. I cannot ensure that this is not a relic of the linear structure (and cannot estimate general a dynamic nonlinear panel data model), and thus do not take this into account in the structural model.

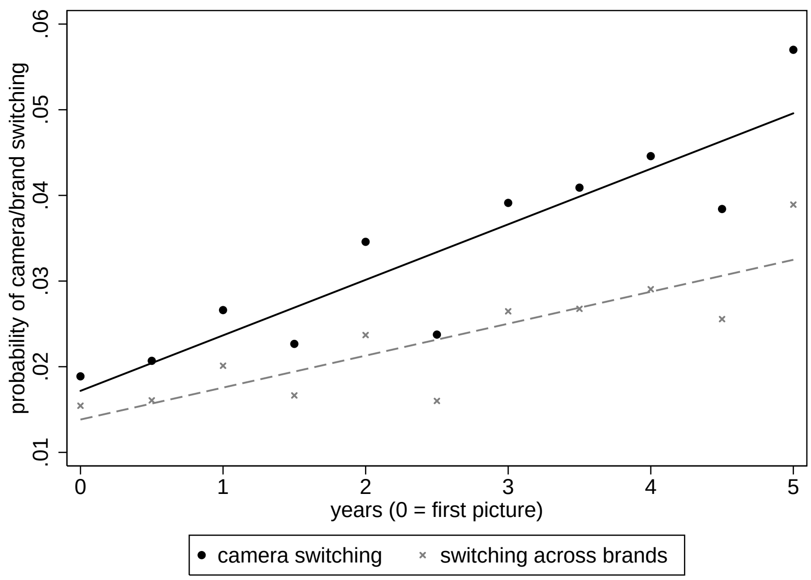

Figure 4: Camera-switching and brand-switching probability

Notes: This figure shows the probability of switching between cameras. To eliminate selective attrition, I limit the sample to consumers we observe for more than 5 years, and focus on the first 4.5 years.

higher experience leads to usage of more advanced cameras. Also shown in the figure, we find that the share of consumers switching across brands are noticeably lower. This is in line with the increasing brand switching costs in their human capital.

5 A structural model

5.1 Overview

This section presents the structural model. Whereas the idea on consumer demand with learning by doing is general, the model will be presented in the context of digital camera markets for concreteness.

In the model, I jointly characterize a consumer's decisions to purchase a digital camera, and her decisions to use the product. Combining a camera and the stock of human capital produces pictures that generate consumption utility, and at the same time, contributes to the consumer's human capital stock. Therefore, past usage decisions build up consumer human capital, and hence future utility. With rational expectations and a non-zero discount factor, the consumer makes camera replacement and usage decisions, taking into account the consequences of her decisions on her future human capital stock.

5.2 Timing

Consumer i in each period $t = 1, ..., T$ decides whether to purchase a new camera and whether to produce pictures. In the process, her camera technology and human capital endogenously evolve as a consequence of her decisions. The timing of her decisions and state variable evolution is as follows: she first chooses whether to purchase a camera, and if purchase, which *format* and *brand* to buy. Given the brand-format combination, she does not know which camera *model* has higher potential quality, and will randomly draw one according to the market distribution at the time of purchase. If she buys a new camera model, she immediately replaces the old one with no resale value. Then, she decides whether or not to take pictures in this period. When she takes pictures, she experiments on a new method – for example, she tries out a new feature on the given camera – and finds the best way she knows to take the picture. Learning occurs if she finds a method that surpasses all the methods she learned in the past, in which case she replaces her existing methods with the new one. Her human capital then evolves. At the end of the period, she derives utility from purchasing a camera, and from the best picture she took, as well as dis-utility from expenditure on the new camera, and effort spent on taking the picture.

I graphically outline the timing assumption in the decision problem, in Figure 5. However, many of the notations are introduced throughout this section. Therefore, the figure might be useful for revisiting the timing assumption.

5.3 Decisions on camera replacement and usage

I denote *all* consumer decisions in a period as $\mathbf{A}_{it} = (B_{it}, D_{it})$, where symbols A, B, and D stand for "action", "buy" and "do", respectively. B_{it} ("buy") characterizes the purchase decision, as a choice over a brand-format combination.[23] I consider two formats (a compact camera or a DSLR),

[23]Because there are over 2000 camera models, I cannot fully realistically characterize the dynamic decision among those within reasonable computation burden.

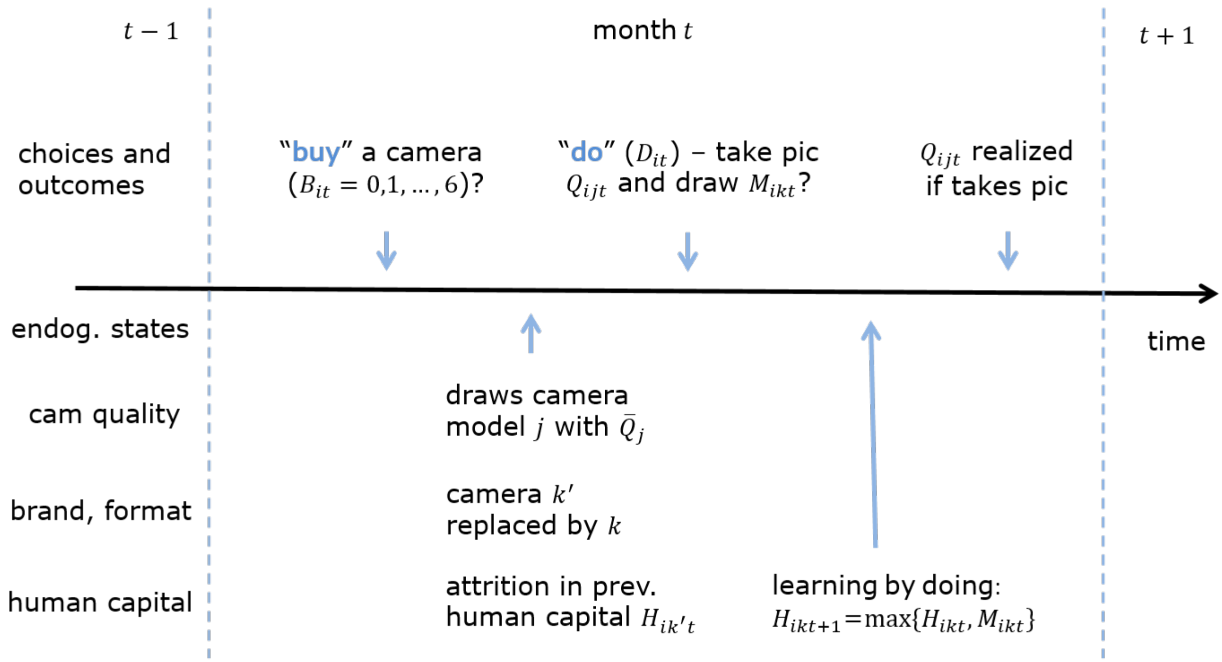

Figure 5: Timing of decisions and evolution of the state variables

Notes: The figure presents the timing assumptions of consumer decisions and state evolution, in a given period.

and three brands (Canon, Nikon and "other brands"), plus the option of not buying ($B_{it}=0$). In the rest of this paper, refer to a brand-format combination as a "camera".

State variable $K_{it}=1,\ldots,6$ denotes the camera, owned by consumer i *at the end of* period t.[24] If a camera is purchased, the consumer replaces the previous camera that she owned with the new one, i.e.

$$K_{it}=\begin{cases} K_{it-1} & \text{if } B_{it}=0 \\ B_{it} & \text{if } B_{it}>0. \end{cases} \tag{4}$$

I do not consider resale, or multiple camera ownership.[25]

Given the choice of brand and format, the consumer knows that the exact camera "model" she will receive has idiosyncratic quality $\bar{Q}_j$, but is not informed of the realization of it before receiving the product. Instead, she expects to draw a realization of $\bar{Q}_j$ from the market distribution at the

[24] I fix $k=1,2,3$ to be the compact cameras, and $4,5,6$ to be DSLRs. Also, I use $\tilde{k}=1,4$ as the realized format, where naturally, 1 refers to a compact camera and 4 refers to a DSLR.

[25] As discussed in Section 3.4, I observe very few cases of possible multiple camera ownership. This does not justify modeling this, as it will greatly complicate computation. Also, I do not observe resell, or whether a consumer purchases from the first or secondary market.

time of purchase, observe it, and keep it until she replaces the product with a new one.

The binary variable D_{it} ("do") denotes the decision of whether to take pictures ($D_{it}=1$) or not ($D_{it}=0$), using the latest camera K_{it} with model j, i.e. after the replacement decision. When taking pictures, she produces picture quality Q_{ijt}, from which she derives her consumption utility. Also, using cameras incurs a cost of effort e_i, a parameter which summarizes the dis-utility from taking pictures in a period.

To keep the model simple, I do not model the decision on the number of pictures to take. This is not central to the mechanism this paper addresses. And modeling the number of pictures will necessitate modeling of picture selection and upload decisions, which – due to that we only observe the selected pictures – is not identified without stronger assumptions.

5.4 Consumer production function

The individual derives utility from the picture quality she produces, which comes from a production function of three arguments: her time invariant characteristics; her human capital stock, and the technology of her camera model. I denote her personal characteristics as a parameter q_i. Her human capital is denoted as H_{ikt}, where we specifically emphasize that the stock of knowledge is with respect to camera k. Finally, the technology of the camera format $\tilde{k}$ is denoted as coefficient $\gamma_{\tilde{k}}$, and the technology specific to a camera model j as $\bar{Q}_j$.

Combining these components, we specify the consumer production function as

$$Q_{ijt}\left(\bar{Q}_j, K_{it}, H_{ikt}\right) = q_i + \bar{Q}_j + \gamma_{\tilde{k}} \cdot H_{ikt} + \eta_{ijt}. \tag{5}$$

where η_{ijt} is an independent and identically distributed (IID) error term, which follows logistic distribution with scale σ_η and location 0. It captures non-systematic variation in the maximum picture quality.

Next, I define consumer human capital and its evolution in Section 5.5, and specify the evolution of technology $\bar{Q}_j$ in Section 5.6.

5.5 The evolution of consumer human capital

5.5.1 Consumer human capital and learning by doing

Taking pictures improves one's ability to take pictures in the future. I model learning by doing as the consumer's ability to replicate good methods – i.e. ways to take good pictures – she used in the past, and apply them in future periods. Specifically, the consumer randomly experiments a method, each time she decides to take pictures. If she finds a method that surpasses all the methods she learned in the past, she replaces her existing method with the new one. Human capital captures a consumer's knowledge on the best method to take pictures. Although good methods that improve human capital arrives stochastically, on average, more experienced consumers are capable of producing better pictures.

To formalize this idea, I assume that potential methods M_{ikt} are normally distributed with individual-specific variance:

$$M_{ikt} \sim \mathcal{N}\left(0, \sigma_i^2\right).$$

If a consumer has no experience, she draws M_{ikt} and uses it (and camera k) to produce a picture. She then keeps the method until she finds a better one. The best method she ever found then defines her human capital H_{ikt}. Similarly, holding certain human capital stock, the consumer learns when a good method is found, i.e. $M_{ikt} > H_{ikt}$. Therefore, learning by doing is the replacement of obsolete human capital:

$$H_{ikt+1} = \begin{cases} M_{ikt} & \text{if } D_t = 1 \text{ and } M_{ikt} > H_{ikt} \\ H_{ikt} & \text{otherwise.} \end{cases} \tag{6}$$

Note that this equation recursively defines consumer human capital.[26]

There are three implications from the characterization of human capital. First, current human capital stock leads to higher *expected* picture quality tomorrow, if camera is unchanged. Second, probability of human capital improvement, $\Pr\left(M_{ikt} > H_{ikt} | H_{ikt}\right)$, is decreasing in the current human capital stock. This implies decreasing learning speed in experience, which is a common feature in many learning models, such as the Bayesian learning model (Erdem and Keane, 1996),

[26] This way of modeling human capital evolution is closely related to Lucas Jr and Moll (2014). In their paper, an individual decides whether to work or learn. In the latter case, she takes a draw from the human capital distribution, and adopts the drawn human capital only if it is higher than her own.

the Bayesian learning by doing model (Jovanovic and Nyarko, 1996), and the characterization of production experience curve (Benkard, 2000 and Besanko et al., 2010). Finally, this is also a stochastic learning model, in line with our descriptive evidence in Section 4.5.

This model of human capital and learning by doing generates three attractive implications. First, higher current human capital stock leads to higher *expected* picture quality tomorrow, using the same camera. Second, if the consumer takes pictures every period, the rate of discovery for better methods, $\Pr(M_{ikt} > H_{ikt}|H_{ikt})$, is decreasing in the human capital stock. This implies decreasing learning speed in experience, which is a common feature in many learning models. For example, the Bayesian learning model in Erdem and Keane (1996), the Bayesian learning by doing model in Jovanovic and Nyarko (1996), and the characterization of production experience curve as in Benkard (2000) and Besanko et al. (2010), all share the same feature.[27] This also corresponds to our reduced form evidence.

As normalization, we impose that the consumer starts with zero human capital:

$$H_{ik1} = 0.$$

This implies that a consumer will always drop all negative draws of method. We also normalize the maximum attainable human capital to 1. This normalization is required because both learning speed, switching cost and the returns to human capital in picture quality are free parameters.

5.5.2 Switching cost

If the consumer decides to switch to camera k from camera k', not all her knowledge about k' is transferable to the new camera. For example, the menu layouts of one camera is different from another, and even if a consumer knows to apply a certain method, the extra time spent figuring out how to change settings might cause her to miss shots.

In Section 4.2, we documented that the picture quality from newly adopted cameras is proportionally lower than those from cameras that a consumer is familiar with. Motivated by this pattern,

[27] In fact, two earlier versions of this paper used, respectively, the Bayesian learning by doing model, and a reduced-form experience curve with decreasing return. The main result stays the same with this version.

I assume that the switching cost is proportional to the current human capital stock:

$$H_{ikt} - H_{ik't} = -s_{ik'k} \cdot H_{ik't}, \tag{7}$$

where, for notation simplicity, we implicitly denoted the *end-of-period* camera brand-format: $K_{it} = k$ and $K_{it-1} = k'$.[28] This structure implies that consumers with a longer history of picture taking have accumulated much experience specific to cameras that they are familiar with, and hence are more locked in to similar products. This is also confirmed in Figure 4.

5.6 Camera technology

The quality of pictures produced by the individual also depends on the technology of her camera. Specifically, camera technology plays two roles in the production function:

First, the *format* of camera, i.e. compact camera or DSLR, complements a consumer's human capital. As shown in the production function, I model the camera format effect to be a parameter on consumer human capital – denoted $\gamma_{\tilde{k}}$ – in her production function.[29] $\gamma_{\tilde{k}}$ is constant across consumers and time, and known to all consumers. This means that with higher human capital, the same improvement in camera technology generates larger changes in picture quality. We provide evidence on this complementary relationship in the appendix.

Second, the characteristics of the camera model j, $\bar{Q}_j$, affects the level of picture quality she produces. Since the researcher is able to observe the productivity level of all products from the data, I model $\bar{Q}_j$ as an index of camera resolution (in integers of mega-pixels), and the year of introduction:

$$\bar{Q}_j = \sum_{r=1}^{35} \psi_{1,r} \mathbf{1}\left(resolution_k = r\right) + \sum_{y=2000}^{2013} \psi_{2,y} \mathbf{1}\left(year_k = y\right). \tag{8}$$

I discuss implementation of this in Section B.2.1.

After the purchase of a new camera, the consumer observes the realization of $\bar{Q}_j$, which stays constant until she replaces the camera. Before purchase, however, she does not have precise knowledge of the technology of a specific camera j, but rather, expects that she will receive the "market

[28] To reduce the number of parameters to be estimated, I restrict that the switching cost is symmetric, i.e. $s_{ikk'} = s_{ik'k}$. In past versions, I allowed for both a proportional and a constant part in the switching cost specification, and find that the main source of switching cost is proportional to human capital.

[29] Recall that we denote camera format $\tilde{k} \in \{1, 4\}$ to denote, respectively any compact camera and any DSLR.

technology level" $\bar{Q}_{m,t}$. In addition, she holds belief that $\bar{Q}_{m,t}$ is first order Markov in

$$\bar{Q}_{m,t} \sim \mathcal{N}\left(\chi_0 + \chi_1 \bar{Q}_{m,t-1}, \sigma_q^2\right). \tag{9}$$

That is to say, before purchase, the expected technology that an individual will get depends on the observed average camera technology among consumers who received new cameras in the previous month.

Similar to Gowrisankaran and Rysman (2012) and Hendel and Nevo (2006), our way of modeling technology index is a dimensionality-reduction assumption. However, different from them,[30] technology in our model is observed, due to the direct measure of consumer production function.

5.7 State space and flow utility

We first clarify the relevant state variables before presenting the utility specification. At the beginning of period t, the consumer decisions depends on the camera she owns at the end of last period, $K_{it-1} = k'$, and the quality of the model j' owned at the end of last period, $\bar{Q}_{j'}$. Her decision also depends on her human capital stock with respect to camera K_{it-1}, denoted $H_{ik't}$. Finally, there are three exogenous state variables: prices for each camera format, $P_{\tilde{k}t}$ for $\tilde{k} = 1,4$; and the market technology level in the previous period, $\bar{Q}_{m,t-1}$. We denote $\mathbf{S}_{it} = \left(K_{it-1}, H_{ik't}, \mathbf{P}_t, \bar{Q}_j, \bar{Q}_{m,t-1}\right)$ for compactness of notation.

We now present flow utility. In the model, the consumer derives per-period utility from from purchasing a camera, and from consuming the picture quality she produces. Also, she derives disutility from the money she spends on the new camera, and from the effort taking the picture. If she does not take pictures or purchase cameras, she derives utility zero plus a random shock. In summary, the flow utility, $\tilde{u}_{it}\left(\mathbf{A}_{it}, \mathbf{S}_{it}\right)$, consists of three parts, plus utility shock:

$$\begin{aligned} u_i\left(\mathbf{A}_{it}, \mathbf{S}_{it}\right) + \varepsilon_{it}\left(\mathbf{A}_{it}\right) &= \left(\alpha_i \cdot \mathbb{E}\left[Q_{ijt} | B_{it}, D_{it}, \mathbf{S}_{it}\right] - e_i\right) \cdot \mathbf{1}\left(D_{it} = 1\right) + \\ &\quad \sum_{k \neq 0} \left(\beta_{i1} P_{\tilde{k}t} + \beta_{i2}\left(P_{\tilde{k}t}\right)^2\right) \cdot \mathbf{1}\left(B_{it} = k\right) + \\ &\quad \sum_{k \neq 0} \sum_{k'} \lambda_{i,k'k} \mathbf{1}\left(B_{it} = k, K_{it-1} = k'\right) + \varepsilon_{it}\left(\mathbf{A}_{it}\right). \end{aligned} \tag{10}$$

[30] Gowrisankaran and Rysman (2012) assume that the discounted sum of future utility is Markov, while Hendel and Nevo (2006) assume that a part of the individual flow utility is Markov.

In the above specification, the first term characterizes the *expected* utility for producing picture quality Q_{ijt}, without knowing its the exact realization. Although we allow the expected utility as a function of all states and actions, it only depends on the brand-format of the previous camera, $K_{it-1} = k'$, purchase decision B_{it}, human capital of the previous camera $H_{ik't}$ – which, together with B_{it}, implies human capital for the current camera – and camera quality $\bar{Q}_j$, which is a function of past camera quality $\bar{Q}_{j'}$ and time t. I allow the marginal utility on picture quality to be heterogeneous across individuals. Denote it α_i. Also, parameter e_i characterizes the effort cost in the attempt to experiment a method and produce the picture(s).

The second term captures the conventional price effects in the consumer purchase decisions. Specifically, I impose quadratic dis-utility from the price spent, when the individual purchases a new compact camera ($\tilde{k} = 1$) or a DSLR ($\tilde{k} = 4$). Because the prices of compact cameras are very different from the prices of DSLRs, one can imagine that the marginal dis-utility from spending an extra dollar might be different on a 80-dollar compact camera, and on a 800-dollar DSLR camera. I allow for a quadratic specification to capture the difference in the marginal dis-utility.[31] To ensure that marginal utility does not change sign within the support of observed prices, I impose a restriction that the turning point of the U-shape does not go below 800 dollars.

Finally, the third term in the utility specification characterizes the immediate (dis-)utility in purchasing a new camera. This include, for example, the psychological effect of choosing a brand that is different from the current camera, or a status effect from purchasing a DSLR (regardless of the quality of pictures one can generate), etc. Further restrictions are placed in Section 5.10.

5.8 Dynamic programming

With rational expectations, the individual makes purchase and usage decisions every period by maximizing the sum of discounted flow utilities, or solving

$$\max_{\mathbf{A}_{i\tau}} \sum_{\tau \geq t} \delta^{\tau - t} \mathbb{E}_t \left[u_i \left(\mathbf{A}_{i\tau}, \mathbf{S}_{i\tau} \right) + \varepsilon_{i\tau} \left(\mathbf{A}_{i\tau} \right) \right].$$

[31] A linear specification will not fundamentally change estimates of the other parameters, but will predict very different elasticities for DSLRs and for compact cameras. On the other hand, a natural log specification will overly flatten the dis-utility profile, within common price range for DSLRs.

Given stationarity assumptions on the function $u_i(\cdot,\cdot)$ (as in (10)) and transition process of $\mathbf{S}_{it} = \left(\bar{Q}_j, K_{it-1}, H_{ik't}, \mathbf{P}_t, t\right)$,[32] this is a standard dynamic decision problem in spirit of Rust (1987) and others, where the consumer solves the equivalent static decision problem

$$\max_{\mathbf{A}_{it}} U_i\left(\mathbf{A}_{it}, \mathbf{S}_{it}\right) + \varepsilon_{it}\left(\mathbf{A}_{it}\right)$$

where the choice-specific value function $U_i\left(\mathbf{A}_{it}, \mathbf{S}_{it}\right)$ is defined by the Bellman equation

$$U_i\left(\mathbf{A}, \mathbf{S}\right) = u_i\left(\mathbf{A}, \mathbf{S}\right) + \delta \cdot \mathbb{E}\left[\max_{\mathbf{A}'} U_i\left(\mathbf{A}', \mathbf{S}'\right) | \mathbf{S}, \mathbf{A}\right]; \tag{11}$$

and all state transition probabilities apply in the expectation operator. I provide detailed calculation of state transition probabilities in the appendix.

5.9 Identification

Section 3.3 discussed identification of implied picture quality, from cross-sectional data of picture taking and posting dates, and their cumulative views. The key identifying assumptions imposed there are that upload date is exogenous given picture taking date, and that the accumulation of potential viewer base (those who might decide to see the pictures, depending on quality, popularity or topics) is determined by calendar time but not individual characteristics. We provide supportive evidence on the first assumption and robustness check for the case when the second assumption fails.

I now discuss parametric identification of the structural model, given (implied) picture quality, choices of picture taking and camera purchase, and other observed state variables as data. First, given a correctly specified model for camera and picture taking choices, the normalization of initial human capital at 0, and some normalization of $\bar{Q}_j$, the production function intercept q_i is identified by the initial period observed picture quality. In fact, I pre-estimate a reduced form model of quality on experience and camera characteristics to obtain the implied $\bar{Q}_j$, so as to guarantee identification of q_i. Next, camera format effect γ_k are identified by comparing differences in the stationary picture quality, across camera formats, because the cap of human capital is normalized at 1. Then, learning

[32] Note that we imposed that the transition of t across two periods can be ignored.

speed is identified by observing changes in picture quality given q_i, $\bar{Q}_j$ and $\gamma_{\tilde{k}}$, before human capital reaches its stationary level. Of course, all of these are conditional on a correctly specified choice model.

For the parameters in the choice model, we first identify parameters in the exogenous state transition matrices, $\Pi_{\tilde{k}}$ as price transition of camera $\tilde{k}$, and parameters that capture market technology index evolution, $\bar{q}_t$ and σ_q. They are identified by the observed prices and technology. We also impose that the discount factor δ is known. It is not identified unless with valid exclusion restrictions (Magnac and Thesmar, 2002). Given these parameters and the production side model, utility parameter α_i is identified by variations in human capital and camera (which changes the expected picture quality), on picture-taking decisions. Effort cost e_i is identified from picture-taking choices when the expected picture quality is zero. Price coefficients are identified by price variations, and other utility coefficients are identified by the "left-over" systematic variations in choices. For example, consumers tend to choose the brand she used before, or they tend to purchase a DSLR even when her human capital does not justify so; these patterns identify, respectively, the brand-switching disutility and the additional utility from purchasing a DSLR camera.

Finally, identification of finite mixture heterogeneity comes from systematic variations in an individual's choices and picture quality outcome. See Kasahara and Shimotsu (2009) for a formal discussion.

5.10 Implementation

5.10.1 Sources of heterogeneity

To capture heterogeneity in the preferences and the human capital formation processes, I assume that there two unobserved, time invariant types of consumers, and observed choices and picture quality are mixtures of type-specific ones. The sources of heterogeneity across individuals could come from differences in preferences, initial human capital, learning speed, and switching costs. This finite mixture setup also allows for arbitrary correlation between the individual-specific parameters. I restrict the production function to be the same across all consumers.

5.10.2 Switching cost

To further parameterize the switching cost $s_{ik'k}$, I allow it to vary across the cases when the consumer switches within the same format of products, or across formats, or across brands. I assume that switching across formats incurs no smaller switching cost than within a format; and similarly, switching across brands incurs no smaller cost than within a brand. To impose these assumptions, I specify the following structure for the switching cost across formats *and* across brands:

$$1 - s_{ik'k} = \left(1 - s_i^{baseline}\right) \cdot \left(1 - s_i^{format}\right) \cdot \left(1 - s_i^{brand}\right)$$

where s_i^{format} and s_i^{brand} symbolize the across-format and across-brand switching cost, taking value 0 when the individual switches within format or brand, respectively.[33]

5.10.3 Choice intercepts and other explanations of state dependence

The utility function in (10) gives a very general specification of choice state dependence and choice-specific intercepts, that does not depend on the potential picture quality one generates. In implementation, I restrict the utility specification to a more parsimonious structure, which is characterized by 5 parameters:

$$\begin{aligned} \sum_{k',k} \lambda_{i,k'k} &= \lambda_{i,DSLR} \mathbf{1}\left(B_{it} \geq 4\right) + \lambda_{i,Canon} \mathbf{1}\left(B_{it} = 1,4\right) + \lambda_{i,Nikon} \mathbf{1}\left(B_{it} = 2,5\right) \\ &\quad + \lambda_{i,FormatSwitch} \mathbf{1}\left(format_{it} \neq format_{it-1}\right) + \lambda_{i,BrandSwitch} \mathbf{1}\left(brand_{it} \neq brand_{it-1}\right) \end{aligned}$$

where λ_{DSLR} captures the immediate utility of purchasing a DSLR camera (relative to a compact camera),[34] $\lambda_{i,Canon}$ and $\lambda_{i,Nikon}$ capture the immediate utility of purchasing specific brands, while $\lambda_{i,FormatSwitch}$ and $\lambda_{i,BrandSwitch}$ capture format- and brand- switching effects (in additional to the switching cost in human capital).

[33] For example, if an individual holds human capital stock of 1 and a Canon compact camera, then, switching to another Canon compact camera costs $1 - \left(1 - s_i^{baseline}\right)$; switching to a Nikon compact camera costs $1 - \left(1 - s_i^{baseline}\right) \cdot \left(1 - s_i^{brand}\right)$; switching to a Canon DSLR costs $1 - \left(1 - s_i^{baseline}\right) \cdot \left(1 - s_i^{format}\right)$; and finally, switching to a Nikon DSLR costs $1 - \left(1 - s_i^{baseline}\right) \cdot \left(1 - s_i^{format}\right) \cdot \left(1 - s_i^{brand}\right)$.

[34] I cannot estimate a separate compact camera utility because the two brand coefficients almost capture the entire market, so a $\lambda_{i,Compact}$ and $\lambda_{i,DSLR}$ together will produce close-to-perfect co-linearity with the brand parameters.

5.10.4 Initial conditions

Heterogeneity in the prior-to-sample experience is characterized by the heterogeneous production function intercept q_i.

Choices of the initial cameras are endogenous to preference, initial quality and learning speed heterogeneity. For example, a consumer with higher learning speed might be more willing to purchase a DSLR camera before period 0. Therefore, her DSLR owned at period 1 is not exogenously given. To endogenize the initial cameras, I compute the stationary distribution of camera formats, conditional on consumer-type specific model parameters and that human capital and market technology are fixed at their initial values. This is similar to Hendel and Nevo (2006).[35]

5.10.5 Discount factor

Finally, I give all consumers a discount factor of 0.99 monthly, implying an annual discount factor of 0.89. Previous versions of this paper has used monthly discount factor of 0.95, and produced similar results.

6 Estimation results

6.1 Camera format coefficients

Table 5 presents parameter estimates for the production function, which is common across different types of consumers. I find that using a DSLR camera improves picture quality, to an extent that will attract 15.5% more views. This improvement is smaller than the difference shown in raw data (e.g. in Figure 15), which also reflects selection of camera formats, by individuals with different skills.

[35] Alternatively, one could model the initial brand-format distributions. I only model the initial camera format distributions because, monthly choice probability being close to zero, the brand-format choice probability matrix is more likely to be singular at some parameter values.

Table 5: Estimates of the production function

	parameter	s.e.
return to human capital - compact camera (γ_1)	0.58	0.07
- DSLR camera (γ_4)	0.67	0.06
scale of quality error term (σ_v)	0.68	0.00

Note: This table reports structural estimates for the parameters that are common across individuals. This includes the share of the first type. Bootstrap standard errors are reported, which are computed from estimates of 20 random samples with replacement.

6.2 Initial quality and learning speed estimates

We then turn to structural parameters that are heterogeneous across individuals, which are all summarized in Table 6.

The constant term in the production function, q_i, is different across the two types of consumers: it is negative for the first type, and large and positive for the second. Note that I normalized initial human capital for all consumers to its lower bound – zero. That means, q_i is the average picture quality in the initial period, net of the camera characteristics index $\bar{Q}_j$ (which is much smaller in magnitude). Hence, we can interpret heterogeneity in q_i as differences in an individual's ability to produce picture quality at the start of the sample. By this result, I label the two types of consumers "low-starters" and "high-starters".

Next, the standard deviation of new method distribution, σ_i, captures learning speed. To see this, recall that an individual can always discard bad draws of methods. Hence, a larger variance implies that she is more likely to draw a better method – which implies higher probability of improving her picture quality each time she takes pictures. With larger σ_i, she also has larger steps of improvement on average. We find that the high-starters have larger σ_i, meaning that they improve faster than the other consumers. To visualize this point, we construct the arrival probability of new methods at each grid point of human capital, shown in Figure 6.

With the observation that learning usually exhibits decreasing return, our learning speed estimates speak in favor of heterogeneity in the consumers' inherent learning rate, rather than differences in their starting position on the same learning curve. High-starters are also fast learners. We visualize the implied distribution of human capital at different cross sections, separately for different types of consumers, in Figure 7.

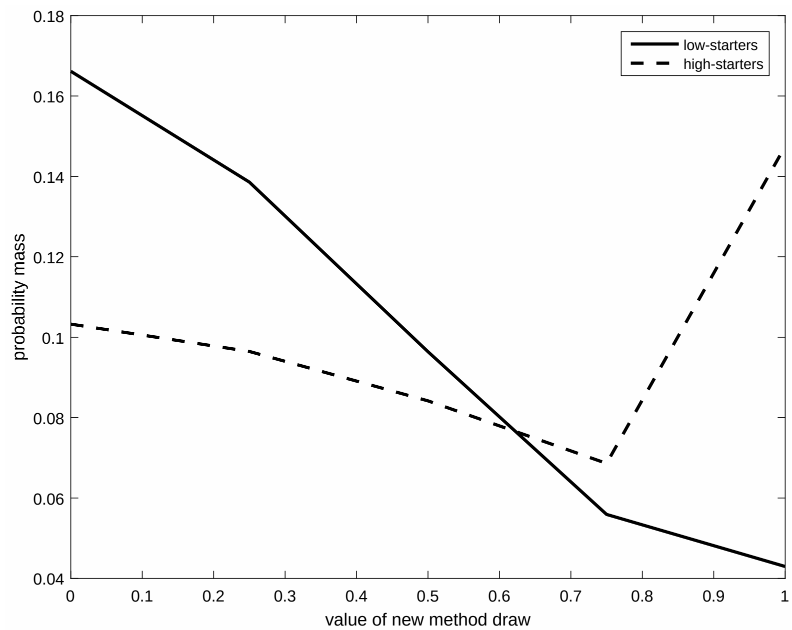

Figure 6: Probability distribution of new methods for different types of consumers

Note: This figure shows the probability mass of new methods arrival, at each grid point of human capital, implied by the parameter estimates. For example, if a high-starter is at zero human capital, there is a 0.1 probability that she draws a new method of 0.5, which will become her next-period human capital. Because negative methods do occur and they will be discarded immediately, the area beneath each probability mass function is 0.5. Note that human capital is constrained in $[0, 1]$, and draws are larger than 1 are capped at 1.

Table 6: Estimates of heterogeneous parameters

	"low-starters"	s.e.	"high-starters"	s.e.
picture quality intercept (q_i)	-0.34	0.06	1.28	0.06
std dev of new methods (σ_i)	0.58	0.09	0.95	0.10
switching cost: baseline ($s_i^{baseline}$)	0.07	0.03	0.11	0.02
- across formats (s_i^{format})	0.08	0.03	0.12	0.02
- across brands (s_i^{brand})	0.11	0.04	0.13	0.02
utility: pref to quality (α_i)	1.80	0.27	1.46	0.26
- effort cost (e_i)	0.23	0.06	1.17	0.49
- price/100 ($\beta_{i,1}$)	-2.38	0.28	-2.69	0.14
- price/100 sqared ($\beta_{i,2}$)	0.13	0.02	0.18	0.01
- preference to DSLR ($\lambda_{i,dslr}$)	6.61	1.05	6.29	0.89
- preference to Canon ($\lambda_{i,canon}$)	-0.04	0.20	-0.14	0.11
- preference to Nikon ($\lambda_{i,nikon}$)	-0.34	0.19	-0.60	0.10
- switching formats ($\lambda_{i,format}$)	-0.22	0.08	-0.35	0.13
- switching brands ($\lambda_{i,brand}$)	-0.32	0.29	-0.91	0.11
type probability	0.61	0.02	0.39	0.02

Note: This table reports structural estimates parameters that are heterogeneous across types. Bootstrap standard errors are reported, which are computed from estimates of 20 random samples with replacement.

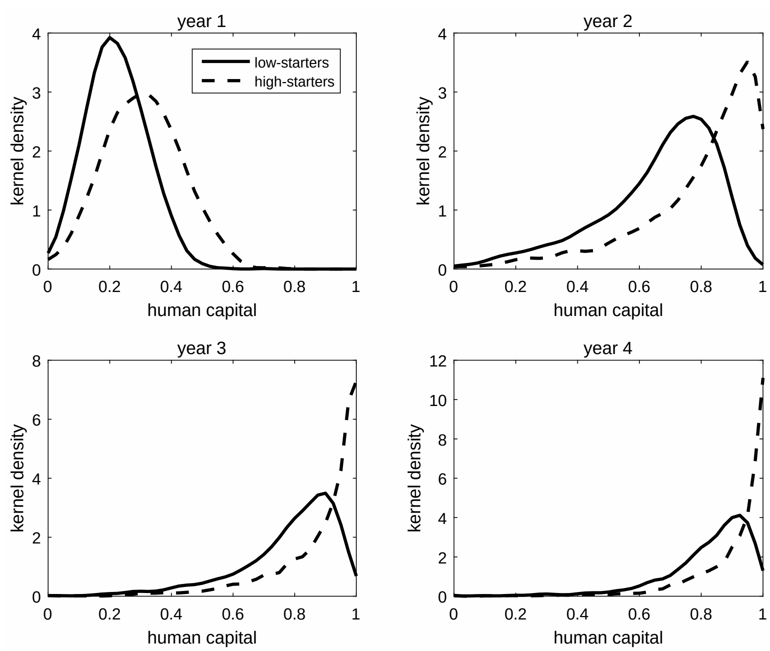

Figure 7: Predicted evolution of human capital

Notes: The four panels present four cross-sections of predicted human capital distributions, separately for different types of consumers.

Table 7: Share of human capital lost upon camera switching

	low-starters	high-starters
same brand/format	7%	11%
across format	14%	21%
across brand	17%	22%
across brand and format	23%	32%

Note: This table presents implied switching cost as a percentage of human capital stock, depending on the type of consumer and the direction of camera switching.

6.3 Switching cost

The switching cost estimates shown confirm our descriptive evidence on the lack of human capital transferability. I find that switching to a camera of the same brand and format incurs some cost beyond the price, in particular for the high-starters. In addition, for all consumers, any across-brand or across-format switch will incur significantly higher loss in human capital.

To better understand the switching cost parameter estimates, I calculate the implied switching cost as a function of the direction of camera switching, and present the numbers in Table 7. For example, for a low-starter with a Canon compact camera, 93% of her human capital is maintained if she switches to another Canon compact camera. On the other hand, switching to a Nikon DSLR will be very costly: it takes away 1/4 of the human capital of a low-starter, and 1/3 of that of a high-starter. We find that camera switching is costly not only in terms of price, but also in terms of the loss of camera-specific knowledge.

6.4 Utility parameters

6.4.1 Preference for picture taking

The estimates on the marginal utility to picture quality, α_i, is insignificantly different between the two types of consumers. This implies that when faced with the same shock in picture quality (from human capital or camera technology), different consumers react in similar ways. On the other hand, we find that there is considerable heterogeneity in the cost of effort, e_i, of taking pictures. Given the expected picture quality, low-starters take pictures much more frequently.[36]

[36] Another way to understand the heterogeneity in e_i, is to think of it as the cost of obtaining human capital. High-starters can take more favorable "lotteries" of new methods, but their "price" for the lottery is more expensive, in terms

6.4.2 Price coefficients

Price effects are very similar across consumers. The nonlinear price effects imply that consumers are more sensitive to 1-dollar price change, at the lower price range. A consumer becomes insensitive to price changes at 800 dollars, which hints that the constraint we imposed on the turning point of the quadratic price effect is binding.

6.4.3 Other utility parameters

The instantaneous utility parameters from camera purchase and brand switching – that are unrelated to picture quality – show that there is considerably positive utility from purchasing a DSLR camera. This might represent the utility from using the advanced features from these cameras, or simply from status effects of using a fancier device. This parameter rationalizes the tendency to upgrade despite at a low human capital level.

For low-starters, the dis-utility from format-switching and brand-switching are insignificantly different from zero. This means that switching cost in human capital captures most of the observed inertia in their brand and camera format choices. For high-starters, they display some inertia when their human capital is at the lower bound. That might suggest that some of their initial skills (reflected in q_i) are brand and format specific.

6.5 Implied price elasticities

We calculate price elasticities, by simulating choice probabilities under the observed prices, and under an instantaneous 5% increase in the price of a given brand-format. The change in price is not expected, and will not persist beyond one month. The own- and cross- price elasticities are shown in Table 8.

I find that the short-run price elasticities are conventional, as in other empirical demand estimation literature in the digital camera industry (Song and Chintagunta, 2003; Gowrisankaran and Rysman, 2012). For example, a 1% decrease in the prices for Canon DSLRs increases the product's current-period demand by 2.6%. This table also shows that most of the additional demand comes from the consumers who would otherwise not purchase in this period (the "no purchase"

of higher effort cost.

Table 8: Average short-run price elasticities

	(1)	(2)	(3)	(4)	(5)	(6)	no purchase
(1) Canon Compact	-2.62	0.03	0.03	0.03	0.03	0.03	0.03
(2) Nikon Compact	0.01	-2.62	0.01	0.01	0.01	0.01	0.01
(3) Other Compact	0.01	0.01	-2.62	0.01	0.01	0.01	0.01
(4) Canon DSLR	0.04	0.05	0.05	-2.99	0.05	0.05	0.04
(5) Nikon DSLR	0.02	0.03	0.03	0.02	-3.18	0.02	0.02
(6) Other DSLR	0.02	0.02	0.02	0.02	0.02	-3.17	0.02

Note: This table reports short-run price elasticities. I compute elasticities by first calculating the implied choice probabilities for each type of consumer, and then the counterfactual choice probabilities when prices for a given brand-format *in a row* are *temporarily* reduced by 5% for the given month. Then, elasticities are computed from the *averaged* choice probabilities. For example, the first row, second column reads: a 1% temporary decrease in the price of Canon compact cameras *decreases* the demand for Nikon compact camera by 0.03%.

category).[37]

I also simulate elasticities to a "permanent" price change, in which case the consumer faces a 5% price change for a brand-format combination, and expects it to last forever. In comparison with the short-run price elasticities, a permanent price change generates interesting demand response, shown in Table 9. For compact camera demand, the own- and cross- elasticities are very similar. For DSLR demand, however, if the price change is committed to be maintained in the future, we find that own-price elasticities are higher than when the price change is one-period only. This is different from typical dynamic demand models – for example in the inventory problems of Hendel and Nevo (2006) – where consumers move future purchases to the current period if they expect the price discount will not last, hence resulting in large short-run elasticities. In addition, if DSLR prices of one brand change permanently, the cross-price elasticities for products in different brand are much higher. This suggests, for example, that consumers who plan to purchase Nikon cameras in the future will systematically avoid doing so, when they expect to switch back due to lower future prices. For the same reason, within-brand cross price elasticities are smaller than when price changes only temporarily, and can sometimes be negative. The permanent price elasticities show that current and future purchase decisions can be inter-temporal complements.

[37] On average, the "no purchase" alternative has a baseline market share of 95%.

Table 9: Average permanent price elasticities

	(1)	(2)	(3)	(4)	(5)	(6)	no purchase
(1) Canon Compact	-2.64	0.03	0.03	0.03	0.03	0.03	0.03
(2) Nikon Compact	0.02	-2.65	0.02	0.02	0.02	0.02	0.02
(3) Other Compact	0.02	0.02	-2.66	0.02	0.02	0.02	0.02
(4) Canon DSLR	0.02	0.42	0.22	-3.07	0.48	0.23	0.04
(5) Nikon DSLR	0.05	-0.24	-0.09	0.15	-3.24	-0.09	0.02
(6) Other DSLR	-0.06	-0.18	-0.17	0.04	-0.16	-2.93	0.02

Note: This table reports permanent price elasticities – computed by changing the price of a brand-format combination by 5% *permanently*.

6.6 Model fit

I also examine whether the model can simultaneously fit the following three patterns: 1) the evolution of a consumer's picture quality along her experience, 2) impact on her picture quality at the instance of camera switching, and 3) her camera purchase and usage probabilities along experience. These are presented graphically in Figure 8.

These figures show that the model is capable of keeping track of both the evolution of picture quality over experience, as well as the evolution of her choice of camera purchase and usage. In addition, the model can capture short-run impacts and recovery of her picture quality, after switching to a different camera.

Note that the model captures causal effects of human capital (and the loss of it upon switching), on picture quality and the incentive to switch between cameras. It also captures heterogeneity in consumers' innate preference and knowledge structure, leading to different choices of camera formats and switching timings. Both the causal effect and heterogeneity is reflected in the observed data patterns: for example, in the bottom panel, the composition of consumers changes in different periods. Therefore, the model fit/prediction should not be read as causal effect.

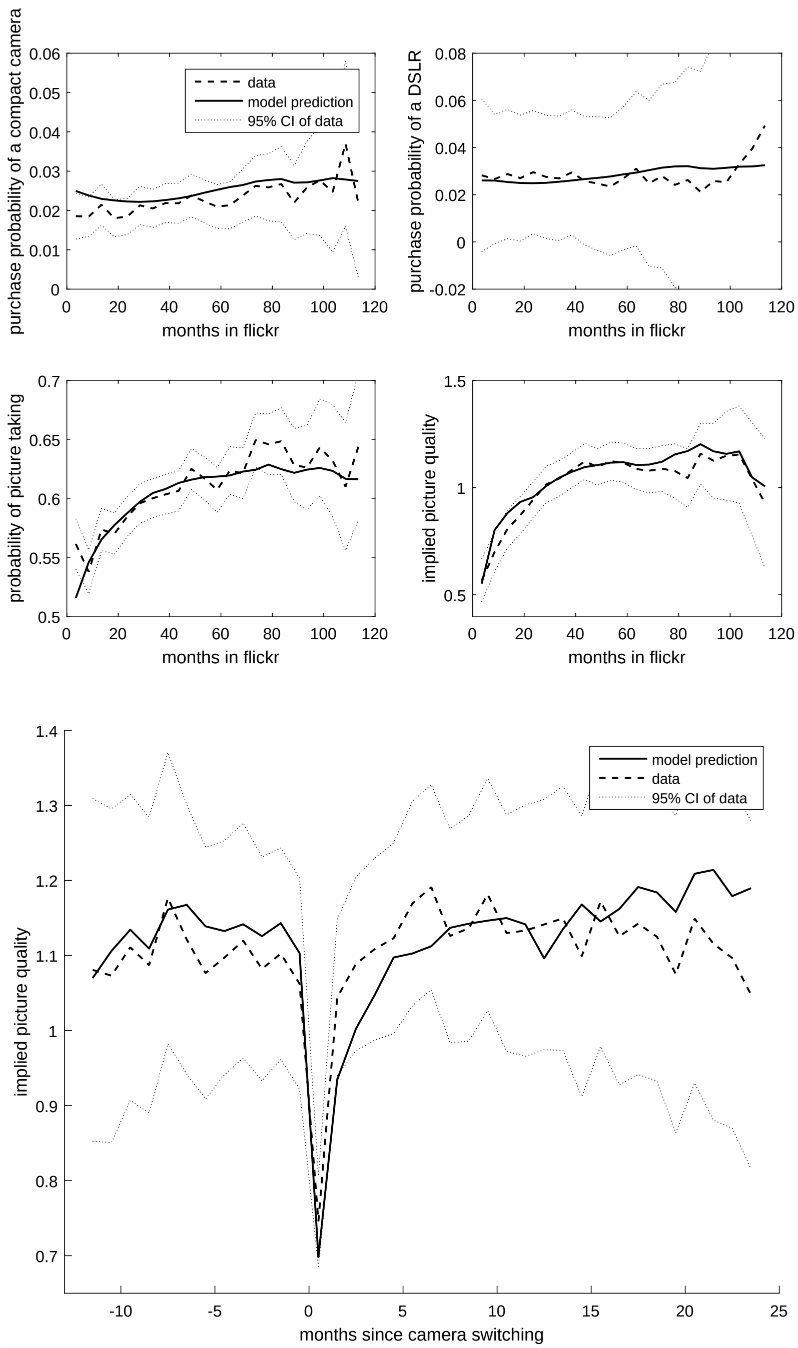

Figure 8: Observed and predicted data

Notes: The top panel presents model predicted and observed data along a consumer's duration in Flickr. The four sub-panels are, respectively, observed and predicted choice probability of compact cameras and DSLRs (upper panels), observed and predicted choice probability for picture taking, and the observed and predicted maximum picture quality in a month. The bottom panel presents the observed and predicted picture quality (measured in log view rate) in a month, by months since camera switching. All predicted values are calculated first conditional on consumer type, then weighted by their posterior type probability.

7 The role of brand-specific learning by doing

7.1 Overview

In this section, I evaluate the impact of consumer switching cost in their brand-specific human capital. Although switching cost is high across camera formats, it is difficult to imagine the counterfactual world where DSLRs and compact cameras are identically designed, or that a user's knowledge can be freely transferred across formats. Brands, however, can easily design their products to be similar; yet they chose not to do so. For example, Nikon DSLR lenses zoom in by turning clock-wise, while Canon lenses turn counter clock-wise. This section simulates counterfactual picture quality, utility and consumer demand for cameras, and evaluates the impact of the brand-specific design "quirks", hindering product usage.

7.2 The size of switching cost

I first simulate counterfactual picture quality when consumer brand switching cost s_i^{brand} is taken away, for consumers who switched between brands.[38] Figure 9 visually show the effect, around the time of camera switching. The solid line is the model-predicted picture quality, holding consumer *choices* and *types* as fixed. Because type distribution is fixed at every point, the drop in picture quality at period 0 measures the causal effect of switching cos.[39] On the other hand, the dashed line is the counterfactual quality under the now-lowered total switching cost.

We find that brand switching cost explains about 1/3 of the total drop in picture quality due to switching cost. A consumer who decides to switch cameras instantly loses 15% audience because of bad picture quality she produces.[40]Without brand switching cost, the loss is reduced to only 10%.

[38]That is to say, consumers who purchase the same camera format but in a different brand now only incurs the base switching cost $s_i^{baseline}$.

[39]Which is smaller in magnitude than observed in the raw data and in Figure 8.

[40]On the other hand, we see that consumers who switch have on average 40% lower view rate right after switching (over all camera choices), and higher loss in view rate if those choices are across brands. The contrast of these two figures show that most of the observed differences in picture view rate before and after switching is due to selection in human capital and camera technology.

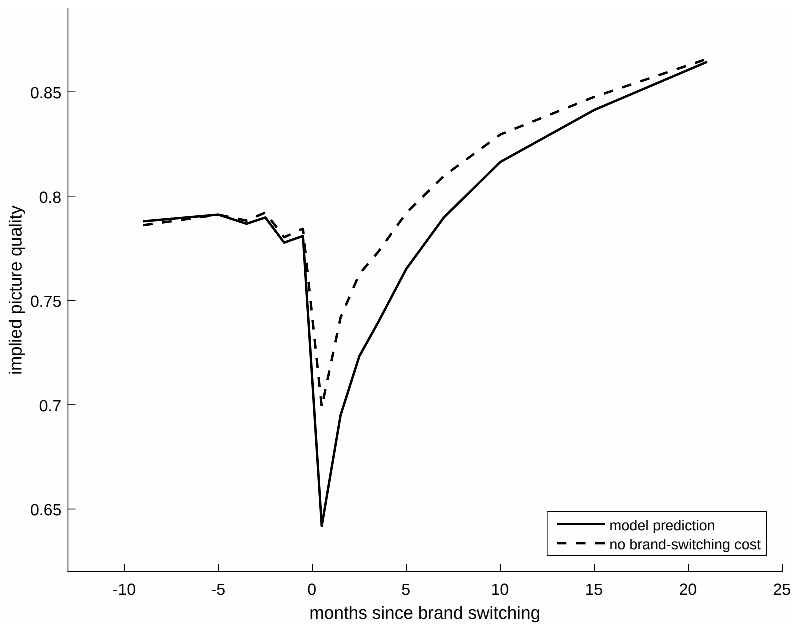

Figure 9: Counterfactual picture quality when additional brand-switching cost is zero

Notes: This figure reports model-predicted picture quality under the parameter estimates (which is the same as in Figure 8), in contrast to when brand-switching cost in human capital is taken away. Specifically, we adjust the counterfactual pre-switching picture quality to the same level as in the model-predicted quality, but contrast their difference after product switching in period 1. Picture quality is measured in log view rate.

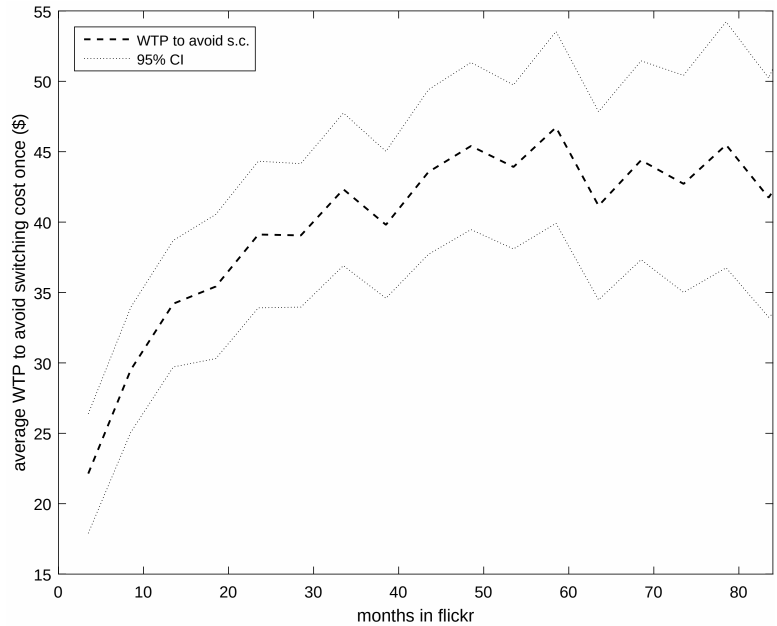

Figure 10: Buyer's willingness to pay to remove switching cost one time

Notes: This figure reports differences in a consumer's discounted sum of utility difference, when she switches brands with and without brand-switching-costs in her human capital. This measures her willingness to pay to remove the brand switching cost one time.

7.3 Willingness to remove switching cost

Given that consumers value picture quality, they will be willing to pay to remove brand switching costs in their human capital. I measure a consumer's willingness to pay for a one-time removal of these costs. This is the discounted sum of utility for the gain in picture quality, with all their knowledge transferred to the new brand, after one purchase. In Figure 10, we plot their willingness to pay, for buyers' with different experience in photography. Intuitively, the willingness to pay measure increases with experience, which on average is increasing in the consumer human capital stock. We also find that a buyer's willingness to pay to remove brand switching cost is high – ranging from \$20 for a beginner to \$40 for an experienced phtographer. In money metric, the brand switching cost is equivalent to considerable price premium for cameras in different brands, and largely discourages brand-switching.

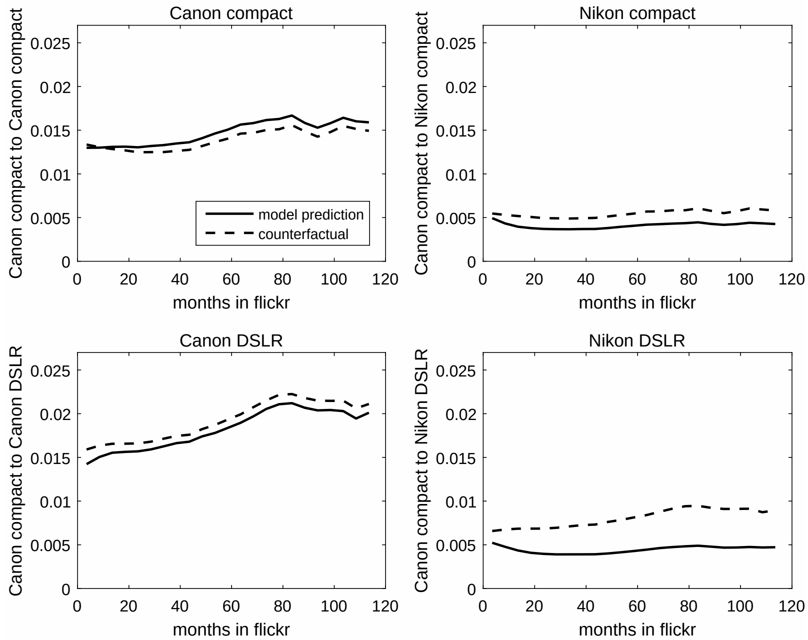

Figure 11: Changes in camera sales for current Canon users, if no brand-switching cost

Notes: Blue solid lines presents the model-predicted probability of purchasing different brand/format of cameras, for consumers currently using Canon. The red dashed lines compare choice probabilities for the same set of consumers, when brand-switching costs are turned off.

7.4 Impact on product demand

I also examine how much the removal of brand switching cost impacts choice probabilities for cameras with different brands and formats. We focus on consumers holding a Canon camera – either a compact camera or a DSLR – so that their demand for Canon and Nikon cameras maps into demand within and across brands. I compare the simulated consumer choice probabilities when the brand switching costs is present, and when it is removed. The average simulated choice probabilities are plotted against experience in photography, presented in Figure 11.

We find that the impact on cross-brand switching (choice of Nikon) is much larger than on within brand re-purchase probabilities, and more so for the experienced consumers. This shows that the elimination of the switching cost relieves the consumers from increasing locked-in to brands, due to the rise of human capital. We also calculate the difference between market shares of Canon and Nikon, among previous Canon users who already makes camera format purchase

Table 10: Relative changes in brand choice gaps, if no brand-switching cost

	compact camera	DSLR
year 1	0.91	0.95
year 2	0.82	0.84
year 3	0.79	0.82
year 4	0.79	0.80
year 5	0.78	0.79
year 6	0.78	0.79

Note: This table reports relative changes in brand choice gaps. For example, the first number reads: the difference in a Canon user's choice probability between buying a Canon and a Nikon, is reduced by 9% if knowledge were freely transferable between brands.

decision.[41] We calculate the relative change of this difference, when shutting down brand switching cost in the consumer human capital. Table 10 reports these results.

Intuitively, difference in brand share measures the persistence in a consumer's choice of brand, which can be due to path dependence or her specific preference towards a brand. Then, relative changes in the difference of brand shares reflects the amount of brand choice persistence captured by consumer switching cost. I find that for experienced consumers, brand-specific human capital accounts for 21-22% of the observed persistence in brand choice. To contrast this, I shut down both the brand-switching cost in consumer human capital, and brand-switching *dis-utilities*, and re-examine the changes in choice persistence across brands. These are presented in Table 4 in the Appendix. Comparing numbers, I find that consumer learning cost attributes to 15-20% of total switching cost of a new consumer; but for an experienced consumer, the loss in brand specific human capital sums up to 1/3 of her total switching costs.

7.5 Impact on cross-price elasticities

Finally, I examine to what extent brand-specific learning by doing drives the substitution pattern across brands. I calculate own- and cross- price elasticities in the same way as in Section 6.5, but taking away the learning cost across brands, s_i^{brand}. Similar to the previous section, I also focus on current Canon users. The price elasticities with and without brand switching cost are shown in Figure 12.

[41] That is, share of Canon given choice of compact (or DSLR), *minus* the share of Nikon given the choice of compact (or DSLR), among previous Canon users.

Most notably, we find that Canon camera sales become about 3 times as sensitive to changes in Nikon prices, when brand switching costs are taken away. This shows that the Canon-specific human capital creates a barrier, that prevents its user base to switch to Nikon. When such barrier is taken away, the same Nikon price drop becomes much more attractive to a Canon user. This implies that across-brand learning cost *protects* Canon's user base from competition, and therefore allows Canon to charge higher markup.

On the other hand, own-price elasticities are not altered by switching costs. Interestingly, the price sensitivity of a Canon user's brand-switching tendency is also not affected by changes in switching costs.

Finally, Figure 12 also documents large differences in price elasticities across different types of consumers. The high-starters, due to their persistently better picture quality (thus higher payoff from using any camera) and larger switching costs, are less elastic to price changes. Because they are more likely to buy DSLR cameras than the low-starters, the price elasticities provide higher markup for DSLRs, as a device to price-discriminate. Price discrimination due to permanent heterogeneity is another important reason that explains the DSLR-compact price gap.

8 Concluding Remarks

The accumulation of consumer's product-specific human capital makes it relatively costly to switch to different products. This makes an experienced consumer more likely to be locked in, and partly explains why her brand choice is persistent in the long run (Bronnenberg et al., 2012). Also, because more experienced consumers value high-end products more, while being less responsive to competitor's price changes, the consumer learning costs explain why high-end products are priced much higher in spite of competitors' presence.

This paper serves as one of the first attempts to measure such consumer learning costs. In the digital camera industry, I measure learning by doing by tracking viewer clicks on pictures, corrected for website-related causes that changes click rates. I find that every time an experienced consumer switches across brands, 10% of her human capital is lost, reflected in the changes in click rates between comparable pictures. For this one-time switching cost, she is willing to pay $40 to avoid. The switching cost explains a quarter of the consumer's persistence in brand choice,

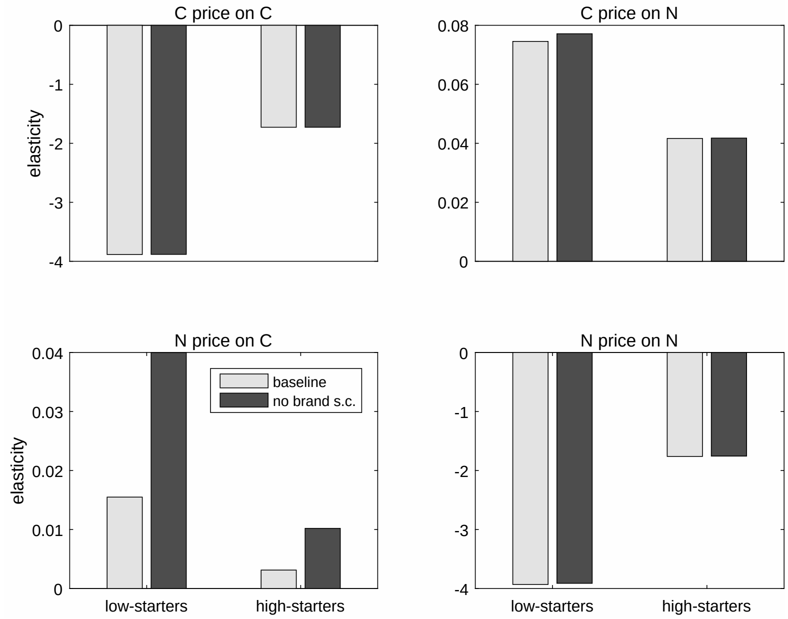

Figure 12: Price elasticities for current Canon users, baseline versus no brand-learning costs

Notes: This figure presents own and cross price elasticities, for current Canon users only, separately for low- and high-starters (left and right groups for each panel), and separately under estimated switching cost and when brand-switching costs s_i^{brand} is removed (light and dark bars). Elasticities calculations are documented in Table 8.

and lowers her cross price elasticity by a factor of 3.

From our estimates, there are several reasons for policies that encourage product design similarities. First, similarity in product design will reduce product-specific learning costs, which prevents wastes in consumer knowledge after a consumer starts using a different brand. Second, *ex ante*, lowering this cost will allow her to switch away from the brands that she does not like, but was locked into due of transitory taste or technology shocks in the past. Finally, this policy may also reduce the prices of DSLR cameras, from which firms currently have excessive market power due to switching costs. However, the reduction of switching cost might lower the incentive to price-penetrate in the entry-level market, so the net welfare effect from price changes is unknown.[42]

The main limitation of this study is that we focus on consumer behavior rather than market equilibrium, and do not estimate aggregate demand using market share data. Computationally, characterizing market equilibrium will be more burdensome (Gowrisankaran and Rysman, 2012)

[42]Klemperer (1995) argue that in most cases, prices will be higher in presence of switching costs. Arie and Grieco (2014) show conditions under which this is true. Dubé et al. (2009) calibrate a model using estimates from consumer package goods, and find that equilibrium prices are lower in presence of switching costs.

and might require some simplification on other parts of the model. Nevertheless, this exercise will allow us to evaluate counterfactual firm pricing policy, and is one important direction for future research.

References

Albuquerque, P. and Y. Nevskaya (2012), 'The impact of innovation on product usage: A dynamic model with progression in content consumption'.

Arellano, Manuel and Stephen Bond (1991), 'Some tests of specification for panel data: Monte carlo evidence and an application to employment equations', *The review of economic studies* **58**(2), 277–297.

Arie, Guy and Paul LE Grieco (2014), 'Who pays for switching costs?', *Quantitative Marketing and Economics* **12**(4), 379–419.

Becker, G.S. (1965), 'A theory of the allocation of time', *The economic journal* **75**(299), 493–517.

Benkard, C Lanier (2000), 'Learning and forgetting: The dynamics of aircraft production', *American Economic Review* pp. 1034–1054.

Besanko, David, Ulrich Doraszelski, Yaroslav Kryukov and Mark Satterthwaite (2010), 'Learning-by-doing, organizational forgetting, and industry dynamics', *Econometrica* **78**(2), 453–508.

Bronnenberg, Bart J, Jean-Pierre H Dubé and Matthew Gentzkow (2012), 'The evolution of brand preferences: Evidence from consumer migration', *American Economic Review* **102**(6), 2472–2508.

Bronnenberg, Bart J, Sanjay K Dhar and Jean-Pierre H Dubé (2009), 'Brand history, geography, and the persistence of brand shares', *Journal of political Economy* **117**(1), 87–115.

Crawford, G.S. and M. Shum (2005), 'Uncertainty and learning in pharmaceutical demand', *Econometrica* **73**(4), 1137–1173.

David, Paul A (1985), 'Clio and the economics of qwerty', *The American economic review* pp. 332–337.

Dubé, J.P., G.J. Hitsch and P.E. Rossi (2009), 'Do switching costs make markets less competitive?', *Journal of Marketing Research* **46**(4), 435–445.

Dubé, J.P., G.J. Hitsch and P.E. Rossi (2010), 'State dependence and alternative explanations for consumer inertia', *The RAND Journal of Economics* **41**(3), 417–445.

Erdem, Tulin and Michael P. Keane (1996), 'Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets', *Marketing Science* **15**(1), 1–20.

Farrell, J. and C. Shapiro (1988), 'Dynamic competition with switching costs', *The RAND Journal of Economics* pp. 123–137.

Gabszewicz, Jean, Lynne Pepall and Jacques-Francois Thisse (1992), 'Sequential entry with brand loyalty caused by consumer learning-by-using', *The Journal of Industrial Economics* pp. 397–416.

Gowrisankaran, Gautam and Marc Rysman (2012), 'Dynamics of consumer demand for new durable goods', *Journal of Political Economy* **120**(6), 1173–1219.

Greenstein, Shane M (1993), 'Did installed base give an incumbent any (measureable) advantages in federal computer procurement?', *The RAND Journal of Economics* pp. 19–39.

Greenstein, Shane M (1997), 'Lock-in and the costs of switching mainframe computer vendors: What do buyers see?', *Industrial and Corporate Change* **6**(2), 247–273.

Heckman, James J (1981), Heterogeneity and state dependence, *in* 'Studies in labor markets', University of Chicago Press, pp. 91–140.

Hendel, Igal and Aviv Nevo (2006), 'Measuring the implications of sales and consumer inventory behavior', *Econometrica* **74**(6), 1637–1673.

Israel, Mark (2005), 'Services as experience goods: An empirical examination of consumer learning in automobile insurance', *The American Economic Review* **95**(5), 1444–1463.

Jovanovic, B. and Y. Nyarko (1996), 'Learning by doing and the choice of technology', *Econometrica: Journal of the Econometric Society* pp. 1299–1310.

Kasahara, H. and K. Shimotsu (2009), 'Nonparametric identification of finite mixture models of dynamic discrete choices', *Econometrica* **77**(1), 135–175.

Keane, M.P. and K.I. Wolpin (1994), ‘The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte carlo evidence’, *The Review of Economics and Statistics* pp. 648–672.

Kellogg, Ryan (2011), ‘Learning by drilling: Interfirm learning and relationship persistence in the texas oilpatch’, *The Quarterly Journal of Economics* p. qjr039.

Klemperer, P. (1987), ‘Markets with consumer switching costs’, *The Quarterly Journal of Economics* **102**(2), 375–394.

Klemperer, P. (1995), ‘Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade’, *The Review of Economic Studies* **62**(4), 515–539.

Levitt, Steven D, John A List and Chad Syverson (2013), ‘Toward an understanding of learning by doing: Evidence from an automobile assembly plant.’, *Journal of Political Economy* **121**(4), 643–681.

Li, Jiaxuan (2015), ‘Gateway products in the dslr camera market: Dynamic demand, consumer learning and switching costs’, *Working paper* .

Lovett, Mitchell J and Richard Staelin (2012), ‘The role of paid and earned media in building entertainment brands: Reminding, informing, and enhancing enjoyment’.

Lucas Jr, Robert E and Benjamin Moll (2014), ‘Knowledge growth and the allocation of time’, *Journal of Political Economy* **122**(1).

Magnac, Thierry and David Thesmar (2002), ‘Identifying dynamic discrete decision processes’, *Econometrica* **70**(2), 801–816.

Melnikov, Oleg (2013), ‘Demand for differentiated durable products: The case of the us computer printer market’, *Economic Inquiry* **51**(2), 1277–1298.

Michael, R.T. (1973), ‘Education in nonmarket production’, *The Journal of Political Economy* pp. 306–327.

Nosal, Kathleen (2012), ‘Estimating switching costs for medicare advantage plans’, *Unpublished manuscript, University of Mannheim* .

Osborne, M. (2007), ‘Consumer learning, switching costs, and heterogeneity: A structural examination’, *Quantitative Marketing and Economics* pp. 1–46.

Ratchford, Brian T (2001), ‘The economics of consumer knowledge’, *Journal of Consumer Research* **27**(4), 397–411.

Rust, J. (1987), ‘Optimal replacement of gmc bus engines: An empirical model of harold zurcher’, *Econometrica: Journal of the Econometric Society* pp. 999–1033.

Shaw, Kathryn and Edward Lazear (2008), ‘Tenure and output’, *Labour Economics* **15**(4), 704–723.

Shcherbakov, O. (2008), ‘Measuring consumer switching costs in the television industry’, *Unpublished manuscript, University of Arizona* .

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.

Appendices

A Returns to experience in photography

A.1 Overview

This section estimates the returns to experience in photography in reduced form, where the model is flexible enough to allow for deviations from the assumptions used to infer picture quality, in Section 3.3. Hence, this section also serves as a robustness check for the picture quality measure.

A.2 Specifications

I first further assume quadratic specification of picture quality q_{ip} from Equation (2):

$$q_{ip} = \theta_1 x_{ip} + \theta_2 x_{ip}^2 + \sum cam_{it} + q_{i0} + \eta_{it} \tag{12}$$

which is a quadratic on experience x_{ip}, plus a set of camera fixed effects, individual fixed effect q_{i0}, and a error term η_{it}. I then regress

$$\log(views_{ip}) = \sum_{t_0,t_1} \Phi_{t_0 t_1} + z_{ip}\psi + \theta_1 x_{ip} + \theta_2 x_{ip}^2 + \sum cam_{it} + q_{i0} + \eta_{it}, \tag{13}$$

and the parameters θ_1 and θ_2 capture the returns to experience. Note that this specification shares essentially the assumption I use to infer picture quality, other than the additional quadratic functional form on experience, and the separability in camera dummies.

There are, however, two potential concerns to the assumptions to Equation (2). First, the flow of viewers could interact with experience, resulting in heterogeneous display-window effect. In other words, $\Phi_{t_0 t_1}$ might be individual specific. The second concern is associated with the timing of upload, i.e. the user might strategically choose the time to upload a picture based on its quality. Both arguments point to the heterogeneity of the display window dummies.

With this in mind, I also estimate the returns to experience on a more-flexible specification. Although this cannot be used to infer picture quality, it serves as an robustness check. Specifically,

Appendix Table 1: Returns to experience in photography

	individual fixed effect	individual-display window fixed effect
experience (100 months)	0.649***	0.547***
	(0.012)	(0.015)
experience sq (0000s)	-0.141***	-0.116***
	(0.005)	(0.007)
camera dummies	Yes	Yes
topic dummies	Yes	Yes
upload order	Yes	Yes
display window	Yes	No
months since joined Flickr	Yes	No
number of pics uploaded	Yes	No
Rsq.	0.110	0.005
obs.	1557214	1560206

Note: This table provides reduced form estimates of the returns to experience in photography. The dependent variable is log of cumulative number of views, per picture level. The first column corresponds to Equation (13), where we infer the returns to experience using within-individual variation, but adding covariates to control for aggregate calendar time trend in Flickr. The second column reports estimates using within-individual-upload-time variations, based on the specification in Equation (14). Measured in *picture quality*, the first specification estimates a 3-year return to experience of 21.3%, or an annualized 6.7%; the second specification estimates a 3-year return of 18.2% – annualized to 5.8%. Within-effect R-squared are provided.

I allow for interactions of individual heterogeneity and the display-window effects, resulting in individual-display-time dummies $\tilde{\Phi}_{i,t_0,t_1}$. Equation (13) now becomes:

$$\log\left(views_{ip}\right) = \tilde{z}_{ip}\psi + \theta_1 x_{ip} + \theta_2 x_{ip}^2 + \sum cam_{it} + \tilde{\Phi}_{i,t_0,t_1} + \tilde{\eta}_{it}, \tag{14}$$

where $\tilde{\Phi}_{i,t_0,t_1}$ now captures a combined effect of baseline picture quality q_{i0} and individual-specific flow of viewers. We can regress (13) controlling for individual-batch fixed effects.

A.3 Estimates

The first column in Table (1) presents the estimation results from Equation (13). I find that the returns to experience is positive within sample period, with a decreasing marginal return. This is consistent with the learning curve measures in Shaw and Lazear (2008), Besanko et al. (2010), Levitt et al. (2013), among others. Quantitatively, the *annual* return to experience in the first 3 years amounts to an increase in picture quality, such that it generates 6.7% more views.

I further check the sensitivity of this result to potential endogeneity problems, as discussed. Shown in Column 2, the robust learning speed estimates are not economically different from the benchmark estimates – hence, the inferred learning curve economically robust to the potential concern of endogeneity.

B Appendix: implementation details of the structural model

B.1 Transition probability of the state variables

B.1.1 Human capital

As explained in Section 5.5, human capital improves when the consumer decides to take pictures in a period, *and* discovers a better method. So if the consumer does not take pictures, human capital stays constant for a period. If she takes pictures, *but* does not discover a better method, human capital also stays equal to the previous period value. Therefore, following Equation (6), given picture taking, the conditional probability density function for the next period human capital to be equal to h, is

$$\begin{aligned}\Pr\left(H_{ikt+1}=h|H_{ikt},D_{it}=1,B_{it}=0\right) &= 0\cdot\mathbf{1}\left(h<H_{ikt}\right)+\\ &\quad \Pr\left(M_{it}\leq h\right)\cdot\mathbf{1}\left(h=H_{ikt}\right)+\\ &\quad \phi\left(h/\sigma_i\right)\cdot\mathbf{1}\left(h>H_{ikt}\right), \end{aligned}\tag{15}$$

where the first term indicates that human capital in $t+1$ cannot go below H_{ikt} if there were no camera switching; the second term indicates that if the method draw was "unlucky", that the consumer did not find a better method than the historical best, human capital will stay at H_{ikt}. The last term captures the distribution of improvement, where ϕ denotes the standard normal probability density function, and one can rewrite this term into $\phi\left(h/\sigma_m|M_{it}>h\right)\cdot\Pr\left(M_{it}>h\right)$ given $h>H_{ikt}$. Here, it is clear that the density of improved human capital depends on whether the consumer could find a better method, and the conditional density of "better methods".

With camera switching from k' to k, human capital first takes a loss due to switching cost, and then undertakes Equation (15). That means, if signal exceeds $\left(1-s_{k'k}\right)H_{ik't}$ – which is the "left-over" human capital after switching – learning will happen. This implies

$$\begin{aligned}\Pr\left(H_{ikt+1}=h|H_{ik't},D_t=1,B_{it}=k\right) &= \Pr\left(M_{it}\leq h\right)\cdot\mathbf{1}\left(h=\left(1-s_{k'k}\right)H_{ik't}\right)+\\ &\quad \phi\left(h/\sigma_i\right)\cdot\mathbf{1}\left(h>\left(1-s_{k'k}\right)H_{ik't}\right). \end{aligned}\tag{16}$$

Note that ϕ – the density of new method arrival – is unchanged. This highlights the assumption that human capital does not alter the underlying method distribution. However, the probability that human capital improves has changed, because the area $h > (1 - s_{k'k}) H_{ik't}$ is now larger. Therefore, compare terms between Equations (15) and (16), although the expected future human capital decreases as a result of switching cost, the expected learning rate increased. In particular, if the switching cost is large, this implication resembles the sharp drop in picture quality after camera switching, but the high learning speed that immediately follows.

B.1.2 Camera

Besides human capital, there are two terms that changes with a camera: the brand-format combination K_{it} (which we refer to as a "camera") and the characteristics-specific quality index $\bar{Q}_j$ (which we refer to as a "model"). We keep track of all possible current and future brand-format combinations. Therefore, the camera evolves deterministically, as characterized in Equation (4).

The quality index stays constant if the camera, say j', does not change. If the consumer switches to a new camera j, as indicated in Equation (8), she draws a new $\bar{Q}_j$ from an exogenously evolving distribution that depends only on calendar time of purchase. This pins down the distribution of $\bar{Q}_j$:

$$\Pr\left(\bar{Q}_j < q | B_{it} \neq 0, t\right) = \Phi\left(\frac{q - \bar{q}_t}{\sigma_q}\right) \tag{17}$$

where $\bar{q}_t$ and σ_q are parameters that are estimated in reduced form, and Φ denotes standard normal CDF.

B.1.3 Prices

Price transition matrices are exogenously given, and only depend on the price of the same format of camera in the current period. That is to say, we allow for two price transition matrices $\Pi_{\tilde{k}}$, where each element of $\pi_{\tilde{k},ij}$ is

$$\pi_{\tilde{k},ij} = \Pr\left(P_{\tilde{k}t+1} = p_j | P_{\tilde{k}t} = p_i\right)$$

where p_i, p_j are discrete grid points of price.

B.1.4 Calendar time

Time is only relevant in characterizing the evolution in the distribution of $\bar{Q}_j$. A fully rational consumer, who takes into account the evolution of time, will have the incentive to wait. This is because the expected camera quality drawn tomorrow will be higher. However, having time in the state variable then generates a non-stationary choice problem, which cannot be solved by iterating on a stationary Bellman Equation.

To solve this, I notice that the evolution of quality index distribution across months is negligible, and therefore ignore the evolution of t in consumer expectations. In fact, if we measure technology as the percentage contribution in the number of views, then *on average*, a camera adopted 1 month later is capable of generating 0.08% more views.[43] However, compared to the change in the $\bar{q}_t$, cameras adopted in the *same month* have a standard deviation of $\sigma_q = 7\%$ in the unit of views. That means, a "lucky draw" of model that is two standard deviation above the mean can generate 14% more views, compared to the "average" cameras adopted in the same period. With this amount of cross-sectional heterogeneity in $\bar{Q}_j$, a tiny shift in the mean (equal to 1% standard deviation) is negligible. Therefore, I safely assume that the consumer takes the distribution of $\bar{Q}_j$ today when forming the expectation for camera quality tomorrow.

Note that this assumption only takes away time evolution between two neighboring months (which is all we need to have a stationary Bellman Equation), but allow time as a relevant state variable in consumer choice probability. For example, our model captures that a consumer's probability of camera switching depends positively on the age of the current camera.

B.2 Other implementation details

B.2.1 Camera quality index

I capture heterogeneity across cameras of the same brand and format by a state variable $\bar{Q}_j$, which, defined in (8), is a function of observed resolution and year of introduction. This implies that the same camera will always have a fixed $\bar{Q}_j$, regardless of who adopts it and when.

From a researcher's point of view, I estimate parameters $\psi_{1,r}$ and $\psi_{2,y}$, so as to infer $\bar{Q}_j$ from observed camera characteristics. Specifically, I estimate a flexible reduced form model of picture

[43] See estimation results in Section 6.

quality, to capture the contribution of different camera characteristics:

$$\begin{aligned} Q_{ijt} \quad = \quad & \sum_{r=1}^{35} \psi_{1,r}\mathbf{1}\left(resolution_k = r\right) + \sum_{y=2000}^{2013} \psi_{2,y}\mathbf{1}\left(year_k = y\right) + \\ & \sum_{e=1}^{60} \psi_{3,e}\mathbf{1}\left(expr_{it} = e\right) + \sum_{\tau=1}^{120} \psi_{4,\tau}\mathbf{1}\left(tenure_{it} = \tau\right) + \\ & \sum_{s=1}^{20} \psi_{5,s}\mathbf{1}\left(cum.switch_{it} = s\right) + \sum_{t_0} \psi_{6,t_0}\mathbf{1}\left(t = t_0\right) + \tilde{Q}_i + \varpi_{ijt}, \end{aligned} \tag{18}$$

by regressing monthly maximum picture quality of individual i using camera model j at time t, against indicator variables of camera resolution and year of introduction, and the cumulative number of months that the individual has taken pictures ($expr_{it}$), the cumulative months since the individual appeared in the data ($tenure_{it}$), the cumulative number of times that the individual switched across cameras ($cum.switch_{it}$), and calendar time and individual fixed effects. I then take projected values of the first two terms to be a proxy of the camera j's contribution to picture quality.

I then simplify the individual's belief about $\bar{Q}_j$, by assuming that $\bar{Q}_j$ is not observed prior to purchase, and the individual only observes the average technology in the previous period, $\bar{Q}_{m,t-1}$, and use it to infer the distribution of technology in period t. To implement this, I estimate a linear model on average technology indices at t, $\bar{Q}_{m,t}$, on previous-period average technology $\bar{Q}_{m,t-1}$. The individual then takes the coefficients χ_0 and χ_1 as known in their belief.

Appendix Table 2: Estimate of market technology index transition

	par. est.	std. err.
constant (χ_0)	-0.004	0.001
technology in the prev period (χ_1)	0.844	0.020
std dev of residuals (σ_q)	0.072	

Note: Estimates of camera quality transition process, outlined in Section 5.6. I first estimate a reduced form of camera contribution in picture quality, as in Equation (18). I then take the predicted index of resolution and year-of-introduction as camera quality $\bar{Q}_j$. Next, I find all $\bar{Q}_j$ at the time of purchase, and estimate its linear specification on average camera technology index in the previous period, $\bar{Q}_{m,t-1}$. This gives coefficients χ_1 and χ_0. Finally, I predict residuals and compute the standard deviation of the error term.

C Other estimation results

C.1 Transition of exogenous state variables

I first estimate how the state variables transition across periods. In structural estimation, these parameters are known to the consumer and the researcher.

First, I present estimates for the "market" camera quality transition. Implementation of this is documented in Section 5.6. We find that the distribution of the market's new camera technology index display strong persistence.

Next, I present non-parametric estimates for the price transition matrices. I discretize prices of DSLRs into grids of $300, and those of compact cameras into grids of $150. Despite being a rather crude discretization, our way of interpolating the value function outside of state space improves precision of the value functions. We find that prices can only stay the same, or move to the grid point immediately next to the current-period price. In addition, at any price level above zero, there is a higher probability for the price to go down one grid, compared to the probability of going up. Together with consumer's rational expectation of this, our model captures a consumer's incentive to wait for price drops.

Appendix Table 3: Estimate of transition probabilities for DSLR prices

	tmr: 0	300	600	900	1200
now: 0	0.999	0.001	0.000	0.000	0.000
300	0.006	0.993	0.001	0.000	0.000
600	0.000	0.013	0.985	0.002	0.000
900	0.000	0.000	0.020	0.976	0.004
1200	0.000	0.000	0.000	0.054	0.946

Note: Estimates of the transition matrix for DSLR prices. This is done before structural estimation, and this matrix is known to all consumers in the model.

D Additional Figures and Tables

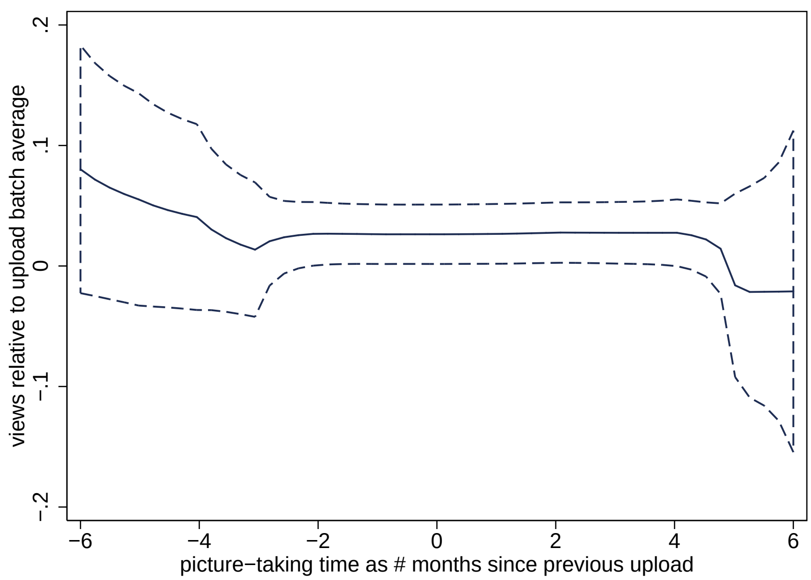

Figure 13: Views on pictures uploaded at different points in time

Note: This figure plots views of a picture relative to views specific to a batch, for pictures uploaded at different points in time. Specifically, the horizontal axis is the *picture taking time*, but is normalized as the number of months since the previous upload batch. For example, -2 means that this picture was taken 2 months before an upload, but the picture was only uploaded in the next batch (so this is delayed upload); +2 means that the picture was taken 2 months after the previous upload and is uploaded in the next batch.

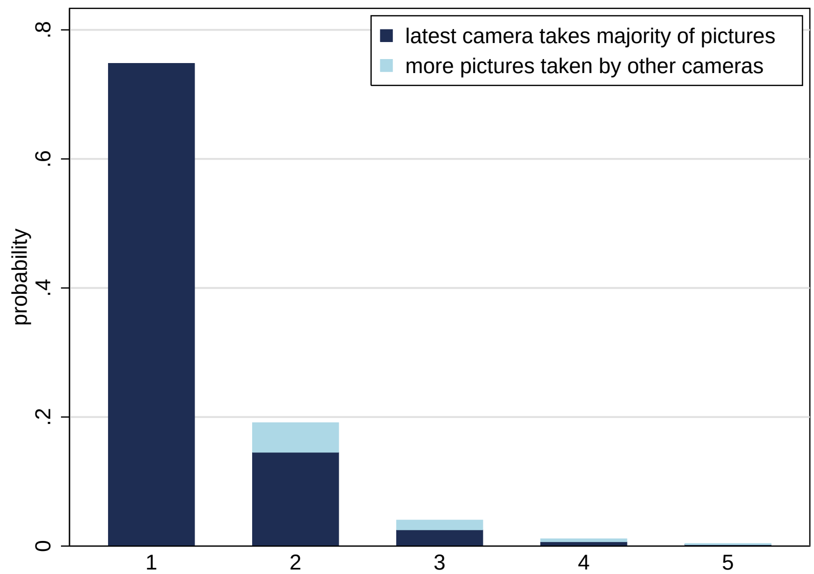

Figure 14: *Joint* probability of the number of cameras owned, *and* whether the latest camera takes the most pictures (x-axis: # cameras)

Note: This figure shows the joint probability of the number of cameras owned at a given time, and the incidence that most pictures are taken by the latest camera. A camera is owned at a point in time if I observe at least one picture taken before that, *and* at least one picture taken afterward. By construction, a camera takes all pictures if it is the only "owned" camera – as represented by the dark bar at x axis = 1. When more than one camera is present, I find that in most occasions, the last camera still takes most pictures.

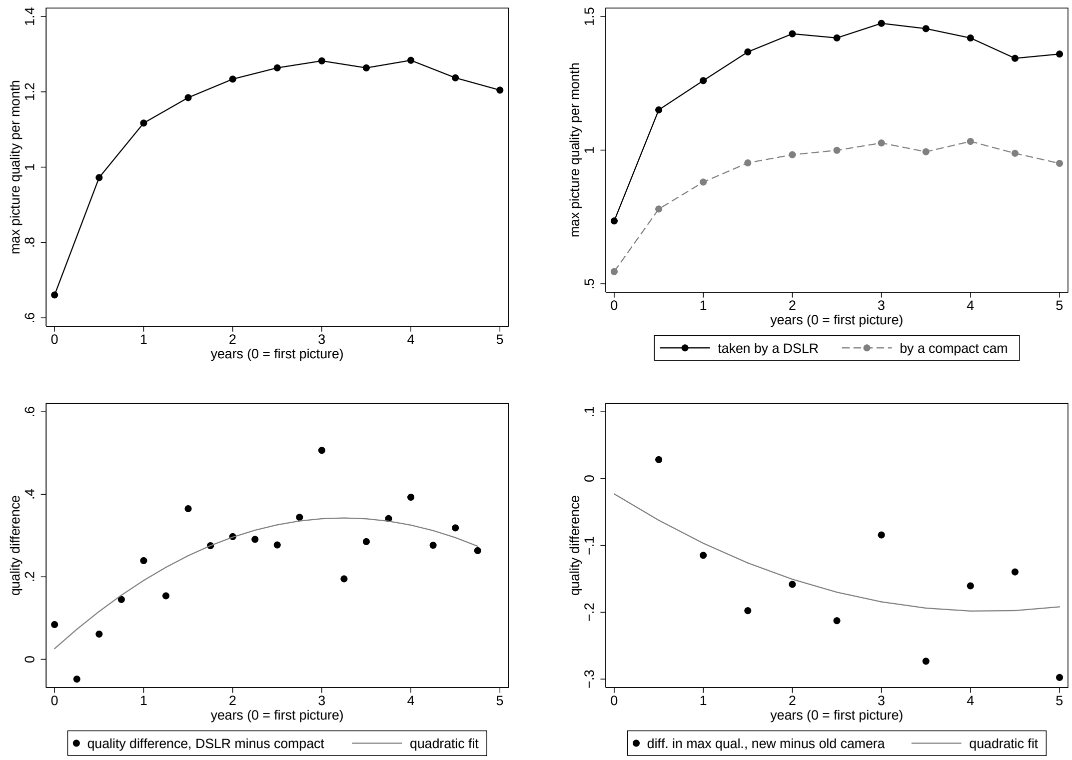

Figure 15: Descriptive figures of picture quality on camera and experience

Notes: These four figures summarize some key aspects of picture quality as a function of experience and camera technology. The upper left panel plots the monthly highest picture quality for a given individual, against experience measured in years. The upper right panel plots maximum picture quality conditional on the format of camera taking the picture, without controlling for the (endogenous) choice of camera format. In the lower left panel, I plot *within-individual* differences in their monthly-maximum picture quality, using the two formats of cameras. That is, we focus on a set of individuals who simultaneously use both formats (within a time window of 6 months), and observe the difference in quality. Finally, the lower right panel plots the difference in the picture quality, for an individual using a camera for the first 3 months (which is likely to be a new camera), and the *same* individual a camera for more than 3 months (which is likely to be the previous camera). The difference then represents (negative of) switching cost, and panel shows switching cost as an increasing function of general experience. The horizontal axis, general experience, is defined as the number of years since one's first picture, *ever* posted on Flickr. To eliminate selective attrition, I limit the sample to consumers we observe for more than 5 years. The markers are mean values across individuals, while the solid line is a sample frequency-weighted quadratic fit. Standard errors are not presented, but most visible differences on the figures are statistically significant.

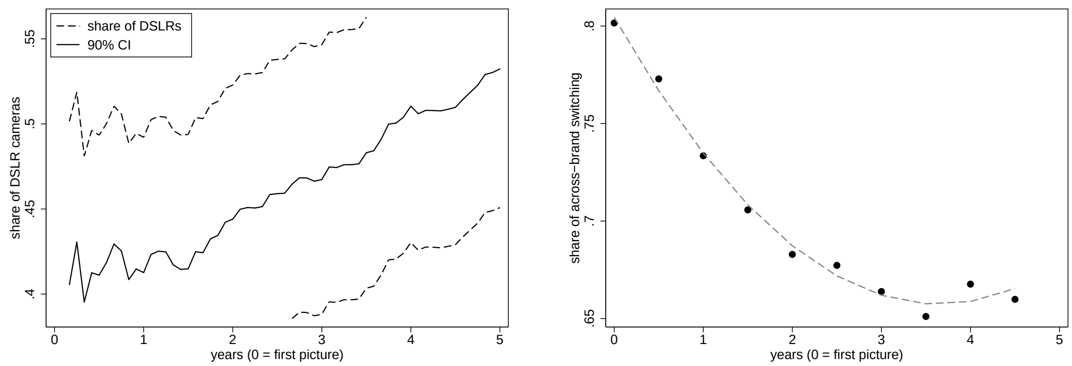

Figure 16: Upgrade and brand-switching probabilities

Notes: The left panel depicts the changes in the choice of product format (DSLR vs compact camera), given a user's years of experience – defined as the number of years since one's first in-sample picture. To control for advances in technology (and other calendar-time effects), I estimate a linear probability model of choice of camera format (DLSR=1), on individual fixed effects, experience fixed effects and calendar time fixed effects. The figure presents the experience fixed effects only. The right panel shows the probability of brand-switching conditional on camera switching. To eliminate selective attrition, I limit the sample to consumers we observe for more than 5 years. So by the end of the figure, the consumers are known to be in the sample.

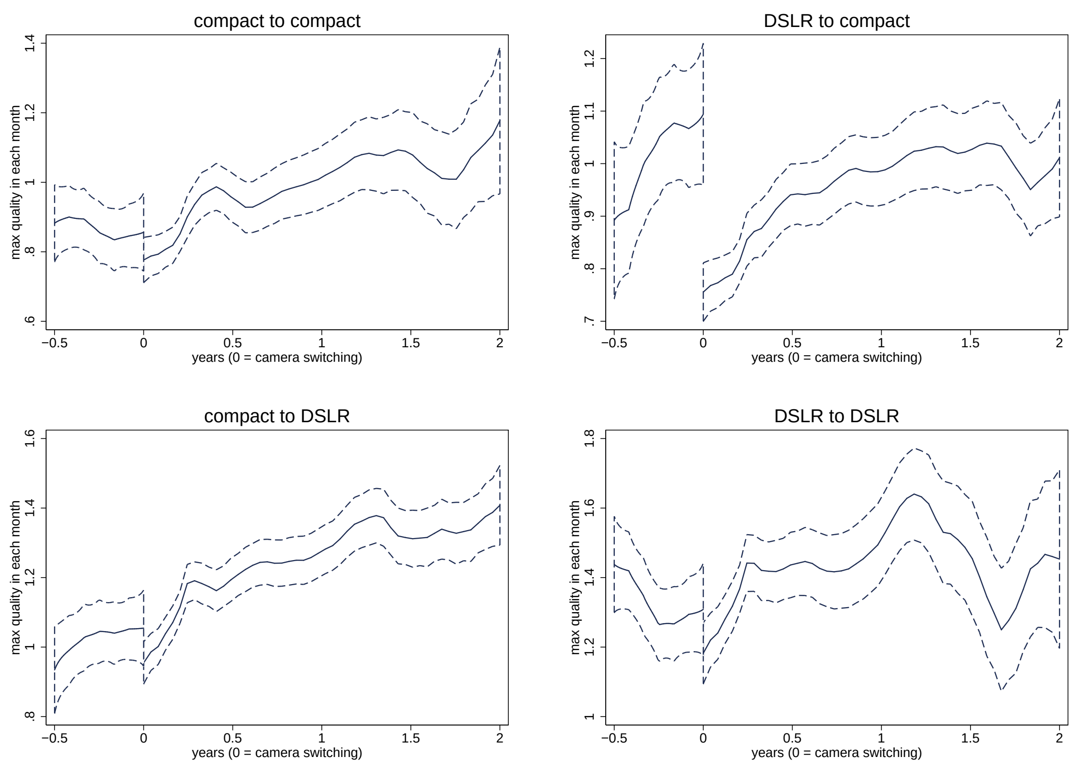

Figure 18: Switching cost controlling for camera formats

Notes: These figures present monthly maximum picture quality for each individual, before and after camera switching, conditional on the camera formats before and after. For detailed notes, see Figure 3.

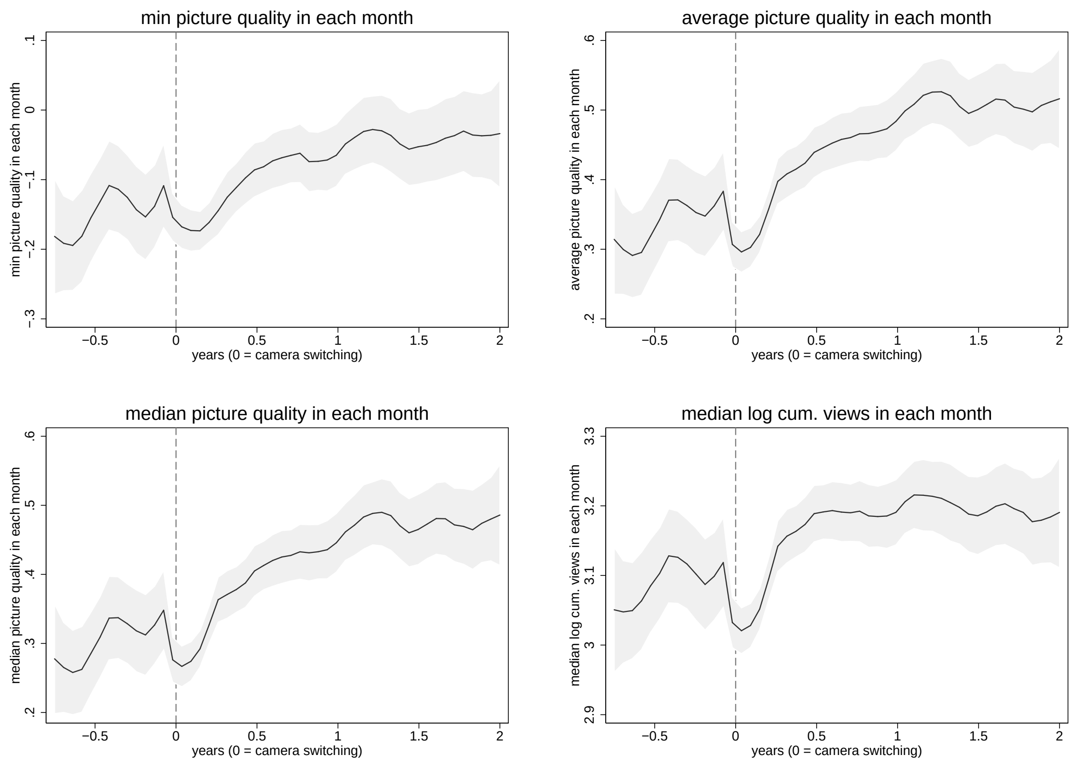

Figure 17: Switching cost under alternative measures of quality

Notes: This is a robustness check of changes in alternative measures of picture quality, around camera switching. The first 3 panels are the minimum, mean and median of the picture quality distribution, around camera switching. The lower right figure presents raw data of log views, i.e. not controlling for display time window effects, topic effects and other Flickr-specific effects. We find that the evidence of switching cost is robust to alternative measures.

Appendix Table 4: Relative changes in brand choice gaps, if no brand-switching cost and dis-utilities

	compact camera	DSLR
year 1	0.46	0.48
year 2	0.41	0.42
year 3	0.40	0.41
year 4	0.39	0.40
year 5	0.39	0.40
year 6	0.38	0.40

Note: This table reports relative changes in brand choice gaps. For example, the first number reads: the difference in a Canon user's choice probability between buying a Canon and a Nikon is reduced by 54%, if knowledge were freely transferable between brands and when brand-switching dis-utilities are eliminated.

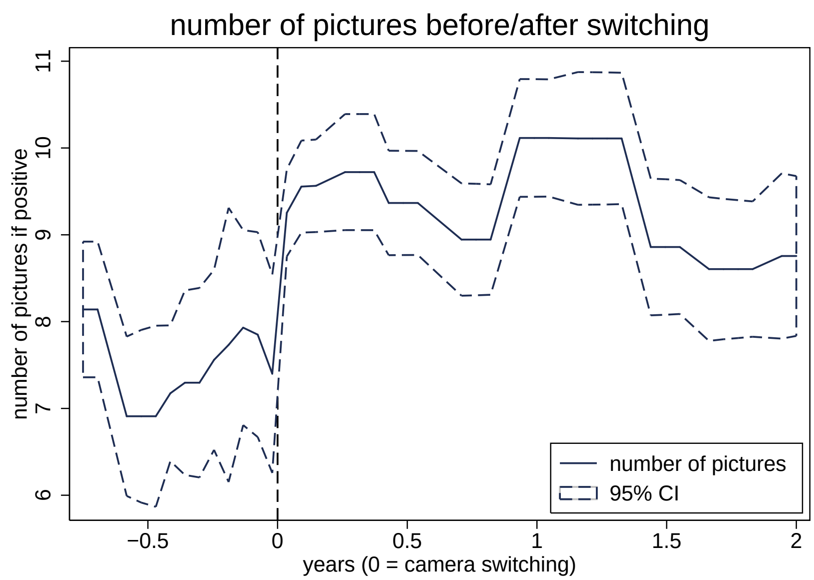

Figure 19: Number of pictures before/after camera switching

Notes: This figure presents the number of pictures an individual produces, around the time of camera switching. This is conditional on these pictures eventually being uploaded to Flickr.