

Demand for Performance Goods: Import Quotas in the Chinese Movie Market

Chun-Yu Ho
University at Albany, SUNY

Marc Rysman
Boston University

Yanfei Wang *
Renmin University of China

September 13, 2020

Abstract

We evaluate Chinese restrictions on the number of foreign movies distributed domestically, particularly an increase in the quota in 2012. We estimate a structural model of consumer demand for movies. We solve a discrete choice model of consumer behavior that is dynamic in the sense that consumers may see movies only once. We find that the reliance on reduced-form age profiles is greatly reduced in our dynamic model relative to standard static approaches. Counterfactual experiments show that consumer welfare increases by 6% due to the import liberalization, and that there is relatively little substitution between foreign and domestic movies.

Keywords: Demand Estimation, Choice Set, Trade Liberalization.

JEL classifications: L10, L82, F13

*We thanks seminar audiences at BU, Carnegie Mellon, Wisconsin, Princeton, Stanford and the NBER.

1 Introduction

Like many developing countries, China restricts the entry of cultural goods such as movies and books. We study the welfare implications of this restriction in the foreign film market from the perspective of consumer choice. We are particularly motivated by China's liberalization of the quota on foreign movies from 20 movies to 34 in early 2012. We ask how much consumer benefit resulted from this expansion, and how much this expansion led to substitution away from other movies, particularly distinguishing between the effect on foreign and domestic movies.

Evaluating welfare from movies is challenging because they are what we call *performance goods*. Performance goods are distinguished by three features. First, performance goods have a frequently evolving choice set. For example, new movies are constantly being introduced, and they typically displace existing movies so that older, but still somewhat recent, movies are often unavailable in theaters for consumers. Second, consumers have limited time to allocate towards consuming movies. Regardless of their income level, consumers would not attend every movie in the theater.

Third, and perhaps most importantly, movies exhibit *consumption durability*. Consumers typically receive significantly lower utility from seeing a movie a second time, so that consumers see most movies only once at most. Consumption durability is a feature of many cultural goods, such as books, museum exhibits, and albums.¹ Many of these goods exhibit stark declines in demand after introduction. Previous research has typically estimated demand for these products with static models that contain an age profile, such as a set of dummy variables for age. Examples are Einav (2007) in movies and Hendricks & Sorensen (2009) for album sales. While this approach may match the data well, it is puzzling from the perspective of economics why the utility from a cultural good would decline at a very rapid rate. A goal of our project is to show that this decline in sales is better explained by a model with consumption durability rather than a reduced-form age profile.

In our model, consumers face an exogenously evolving choice set. Consumers have heterogeneous preferences over movie characteristics, which do not change over time. We assume consumers can see no more than one movie per week, reflecting consumers' limited time for attending cinemas.

¹Consumption durability has long been considered in the literatures on macroeconomic and finance to understand consumption dynamics (Hayashi, 1985; Ferson & Constantinides, 1991).

Further, we assume that consumers cannot see a movie more than once. Thus, the choice set of a given consumer evolves endogenously as the consumer makes decisions over which movies to see. In estimation, we find the level of unobserved quality for each movie-week that rationalizes the observed market share and form a GMM estimator around this term. For much of the paper, we assume that consumers choose myopically which movie to see. Under this assumption, consumers do not account for how seeing a movie today affects future outcomes. We also consider a model of perfect foresight but we show that forward-looking behavior does not fit our movie data well.

We apply our model to a data set covering national box office revenues by week from Chinese movie theaters from January 2012 to June 2015. We collect movie characteristics, such as whether the movie is foreign or domestic, the genre of the movie, and the run-time. We augment the data with a survey from a consulting firm that reports how often people go to the movies. Forcing our model to match this “micro-moment” significantly impacts the results.

Although the drop-off in sales that movies experience from week-to-week is extreme, our results show that it can be entirely explained by consumption durability. In particular, we estimate a traditional static random coefficients logit model with a reduced-form age profile and find that the age profile is strongly significant and negative, reflecting the steep dropoff in sales over the life of a movie. However, estimating our dynamic model with the age profile coefficients reduces the importance of the age profile, and when we impose the micro-moment, we find that the coefficient on the age profile is insignificantly different from zero and precisely estimated. We also find substantial heterogeneity in preferences for foreign movies, suggesting that foreign and domestic movies are not close substitutes.

Because the liberalization going from 20 to 34 movies takes place just before the start of our data, we cannot evaluate the market before the policy change. Rather, we employ our structural model to determine outcomes in the counterfactual scenario. In order to determine which of the 34 movies get dropped when switching to 20 movies, we estimate a probit model of the decision-making by the Chinese government over which movies will be chosen. We find that box office revenue is an important determinant, but the government also considers other criteria such as the rating (PG, R)

and the nationality of the producing firms.

We use the results of this estimation to generate a distribution of possible movies in China if there was no liberalization. We show that consumer welfare increases by 5.98% due to the liberalization. However, the welfare effects of producers are heterogeneous. The import liberalization reduces the total market share of the competing foreign movies more than domestic movies because the extra foreign movies are closer substitutes. This result raises questions for the value of infant industry policies, as substitution between the foreign and domestic products is limited. In addition, we find that if the consumption durability in preferences is ignored, the welfare benefit for consumers is overestimated and the difference in the business stealing effect of foreign movies on competing foreign movies and domestic movies is also overestimated.

Countries may restrict the entry of cultural goods in order to protect domestic industries and also to protect the distinctive nature of their culture from global incursion. We evaluate the implications of the quota only for economic outcomes, such as consumer welfare. Thus, for a policy-maker considering such cultural or industry protection, we provide a measure of the economic cost. Note that in our counterfactual calculations, we assume the set of movies does not change. However, some research and popular press argue that Chinese policies in particular affect movie production in terms of genre and content (see for instance Leung & Qi, 2020). We do not address that issue here, although that it may be important.

2 Literature

Our work contributes to a growing empirical literature on trade in motion pictures. Marvasti & Canterbury (2005) construct a trade barrier index for 33 countries and find that their trade barrier index is positively correlated with imports of U.S. motion pictures. Hanson & Xiang (2011) develop a model of trade with heterogeneous firms for the motion picture industry. They find that average revenues per U.S. film vary widely across countries and are negatively correlated with geographic distance, linguistic distance, and other measures of trade barriers. Thus, these two papers find mixed results of trade barriers for imports of U.S. movies. Holloway (2014) examines 1,236 U.S. movies

released between 1995 and 2004, and finds that movies with a higher quality, measured by their box office in the United States, are more likely to enter into foreign countries. McCalman (2004) studies the role of protection of property rights in the international distribution of movies. Ferreira, Petrin & Waldfogel (2016) estimate a structural model to evaluate the role of product quality in determining gains from trade in motion pictures. Our work differs from those studies in that it uses a structural demand model to examine the welfare effects of import liberalization of U.S. movies.

Our paper builds on the methodology developed by Berry, Levinsohn & Pakes (1995) to estimate a demand system for differentiated products with market-level data. Our work also contributes to three strands of literature related to demand estimation based on Berry et al. (1995). First, we add to the empirical literature on demand estimation for movies. Davis (2006) and Sunada (2012) estimate the effect of the spatial location of a theatre on movie demand. Einav (2007) estimates the seasonality of movie demand. Moul (2007) estimates the effect of word-of-mouth on movie demand. Moul (2008) estimates the conduct of the distributor on rental pricing and advertising. de Roos & McKenzie (2014) estimate the price elasticity of movie demand by exploiting the ticket discount offered by Australian theatres on Tuesdays.

Second, we add to the literature evaluating the welfare benefits of new goods with discrete choice demand models (Trajtenberg, 1989; Petrin, 2002). There are recent studies extending demand models to accommodate some features of cultural goods, such as complementarity between existing offline and new online versions of the product (Gentzkow, 2007) and the unpredictable product quality of new products (Aguiar & Waldfogel, 2018).

Third, we add to the literature on modeling heterogeneous choice sets across consumers in demand estimation. The existing literature suggests that there are two main reasons for having heterogeneous choice sets across consumers. First, the choice sets vary across consumers because some products stock out when consumers make purchase decisions. Conlon & Mortimer (2013) use an expectation-maximization (EM) algorithm to account for the missing data on product availability faced by each customer. Musalem, Olivares, Bradlow, Terwiesch & Corsten (2010) employ a Bayesian method to impute the entire sequence of sales to model product availability faced by each consumer. Second, the

choice sets vary across consumers because of the awareness of different brands. Goeree (2008) models the probability that a consumer would be aware of a given brand. Honka, Hortaçsu & Vitorino (2017) study demand for banks with survey data on consumer consideration sets. Draganska & Klapper (2011) also incorporate information on choice sets from a consumer survey. Barroso & Llobet (2012) model the probability of consumer awareness of a brand as a function of the history of advertising expenditures. Bruno & Vilcassim (2008) show that demand estimates are biased if varying product availability across consumers is ignored.

In addition, we provide a model of consumption durability, in which consumer demand is a dynamic process. Our model is designed for aggregate data (that is, product-level market shares rather than individual choices) and our solution method is similar to Gowrisankaran & Rysman (2012). Relative to that paper, we focus on the durable nature of choice rather than forward-looking behavior, and the formation of choice sets is quite different. While we focus on a model of myopia, we also estimate a model of perfect foresight and estimate the discount rate in the spirit of Magnac & Thesmar (2002). Other papers that estimate the discount rate are Lee (2013), Dalton, Gowrisankaran & Town (2019) and De Groote & Verboven (2019).

Our paper is related to the lengthy literature on the benefits of greater product variety in international trade.² Some observers argue that cultural goods and services “encompass values, identity and meanings that go beyond their strictly commercial value” and request exceptions in protecting domestic cultural goods and services. A separate reason to restrict the entry of foreign goods is to support local producers, so-called infant industry protection. See Greenwald & Stiglitz (2006) and the literature that follows. Our result about limited substitute between foreign and domestic movies suggests this concern is of limited importance.³ For example, Article IV of the GATT agreements in 1947 provides the conditions under which countries may impose quotas on foreign movies.⁴ The

²A leading example is Krugman (1979). The welfare gain from more product variety from trade appears quantitatively large for manufacturing sectors. See Feenstra (1994), Broda & Weinstein (2006), Blonigen & Soderbery (2010) and Sheu (2014).

³Francois & van Ypersele (2002) and Rauch & Trindade (2009) argue that restrictions on trade in cultural goods can raise welfare. Chu-Shore (2010) reports that there is a homogenization of cultural goods in response to trade liberalization. Maystre, Olivier, Thoenig & Verdier (2014) provide theory and evidence to support that trade integration leads to convergence in cultural values across countries.

⁴Many countries impose trade barriers on foreign movies. (Marvasti & Canterbury, 2005) shows that non-tariff trade barriers, such as quotas, are more commonly imposed than tariffs, especially for developing countries.

protection of national culture also played a role in the Uruguay Round of the GATS, which ended in 1994, and the UNESCO Convention on the Protection and Promotion of the Diversity of Cultural Expressions (in particular Articles 6 and 8). There is particular concern for U.S. movies as U.S. producers increasingly rely on foreign revenues and U.S. movies dominate the market share in many foreign countries.⁵

3 Institutional Background

This section discusses Chinese import policies for foreign movies. Until 1994, foreign movies were purchased mainly on a flat-fee basis. Between 1978-1993, the China Film Group was the only authorized agent to import and distribute these films. In each year, the China Film Group spent about USD \$1 million to import about 30 foreign movies, and each foreign movie was purchased at about USD \$30,000. As a result, the imported movies were usually considered “outdated and low-grade but cheap.”⁶

In 1994, the Film Administrative Bureau, under the Ministry of Radio, Film and Television adopted a revenue-sharing practice to import 10 foreign movies per year. The policy aimed to stimulate declining movie attendance and create opportunities for domestic studios. China was approved to join the WTO in 2001. Under the agreement, China increased the quota for revenue-sharing movies to 20. In order to diversify the imported films, in 2004, the State Administration of Radio, Film and Television (SARFT) reserved about six slots for non-U.S. movies.

China has become the largest foreign market for U.S. movies as the annual box office in China has been accelerating faster than 20% annually during the past decade. Specifically, the box office of U.S. movies in China was at \$USD 4.8 billion in 2014. In February 2012, China agreed to significantly increase market access for U.S. movies in order to resolve a WTO dispute that the United States

⁵Marvasti & Canterbury (2005) report that export revenues are an increasing portion of total revenue for U.S. movies. Export revenues were less than one-third of domestic box office revenues in 1986, but were about 90% of domestic box office revenues in 2000. Hanson & Xiang (2009) document that U.S. movies acquire more than 70% of box office revenue in 19 European countries over the period 1995-2004. According to a report by the Motion Picture Association of America, the global box office for U.S. movies released in each country around the world reached \$USD 36.4 billion in 2014, of which, \$USD 26.0 billion was acquired from the international box office. Source: <http://www.mpa.org/wp-content/uploads/2015/03/MPAA-Theatrical-Market-Statistics-2014.pdf>

⁶Stanley Rosen, “The Wolf at the Door: Hollywood and the Film Market in China,” in *Southern California and the World*, eds. Eric J. Heikkila and Rafael Pizarro (Westport, CT: Praeger, 2002), 49–77.

had filed in 2007. With immediate effect, China enlarged its quota for revenue sharing imports of foreign films from 20 to 34 per year. The extra 14 films were specified to be in 3D or IMAX formats. In addition, revenue sharing was set at 25% of box office revenues instead of the previous rate of 13-17%. All of the 34 revenue-sharing movies and all movies imported under the fixed fee plan are imported and distributed by the China Film Group, and some are co-distributed by Huaxia, which is a state-owned enterprise established in 2003. There is no specific quota to import movies on a flat-fee basis, and it is usually 20-30 per year. A small number of top movies have entered by fixed fee in recent years.

A third option for movies to be distributed in China is for them to be co-produced. In a co-production agreement, a foreign producer collaborates with a Chinese investor. In addition, the movie must be sufficiently oriented towards the Chinese market, which SARFT interprets to mean that the movie must feature Chinese actors, Chinese settings, and Chinese themes. Foreign producers obtain attractive revenue-sharing terms (45%) and are not subject to the quota. A challenge is that producers cannot be sure of their co-production status until SARFT reviews the movie. A well-known example in China is *Ironman 3*, which was planned as a co-produced movie but was turned down by SARFT as not being sufficiently Chinese after it was produced. The movie entered China under the fixed fee plan. The movie *Looper* had a similar experience. A successfully co-produced movie was *The Great Wall*. There were only 14 co-production movies from the United States over 2001-2016 (Kokas, 2017, Appendix 1).

All films, foreign and domestic, face censorship by SARFT. Foreign films face censorship regardless of whether they are under a fixed fee plan, under revenue sharing, or are co-produced. Review usually takes 30 days. Article 25 of the Regulation on the Administration of Movies effective in February 2002 prohibits ten aspects of content that would not be allowed in any imported films. The list includes, among other things, “endangers the unity of the nation, sovereignty or territorial integrity,” “propagating evil cult or superstition,” and “propagating obscenity, gambling, violence, or instigates crimes.”

Figure 1 depicts that the share of domestic movies at the box office remained at about 55% over a




Figure 1: Chinese Box Office Revenue and Domestic Share.

long period, which is higher than those in European countries documented in Hanson & Xiang (2009) and may relate to the import restriction of China on foreign movies.⁷ Interestingly, the domestic share does not appear to change much as a result of the liberalization in 2012. As discussed in the next section, we do not rely on pre-2012 data in the rest of the paper, as we view it as less reliable. However, this result foreshadows our finding that there is significant differentiation between foreign and domestic movies.

Similar to other markets, price variation in the Chinese movie market is limited. While prices vary by time of day, day of week, and theater within a cinema, they do not tend to vary by movie, and which movie appears in which theater is a choice made by the cinema.⁸ Similar to Einav (2007) and others in this literature, we do not attempt to estimate a price coefficient. We will capture the mean level of movie utility with a movie fixed effect, and present counterfactual results as percentage changes from the observed outcome. We compute a dollar value as a “back-of-the-envelope” calculation rather than as part of our formal estimation.

⁷The data for this figure were collected by the authors from several on-line sources, particularly reports by Entgroup.

⁸Orbach & Einav (2007) discuss this issue in the U.S., which exhibits constant prices not only across movies but also across time of day and week.

4 Data

The empirical analysis is based on a novel dataset from SARFT of China. The data contain information on box office revenue, the number of tickets sold, and the number of showing screens of all movies shown in each week. Beginning in January 2012, SARFT implemented a system in which cinemas participated in an electronic ticketing program, which greatly enhanced the accuracy with which SARFT could measure these variables. Our data is drawn from SARFT's program. Our empirical analysis includes the movies with admission share for the week larger than 0.1% in at least one week from January 2012 to June 2015. There are 939 movies shown in 183 weeks. We supplement this dataset with hand-collected information on movies, such as genre, run-time, the release date, whether a movie is in 3D or IMAX format, and the nationality of the producing firms.

Table 1 presents a description of the characteristics that we use in our paper. The table presents simple means of the variables, as well as means weighted by ticket sales. The table also breaks out the variables by foreign and domestic movies. We see that foreign movies are more likely to be 3D, IMAX, and action movies, especially when weighted by ticket sales. For instance, 12% of domestic movies are produced in 3D, whereas 44% of foreign movies are produced in 3D, which represents 71% of foreign ticket sales. Similarly, 29% percent of foreign movies are in IMAX relative to 3% of domestic movies, and foreign IMAX movies represent 70% of foreign ticket sales. Foreign movies are more likely to be action movies and less likely to be comedies or dramas, and this is even more extreme when we weight by admissions.⁹

As is common for cultural goods such as books and music, market shares for movies are highly skewed. For each week in our sample, we calculate the share going to each rank of movie, i.e. the top-ranked movie, the second-ranked movie, and so on. We average this over the 181 weeks in our data, and graph the results in Figure 2. The top-ranked movie at 38% is more than 70% higher than the second-ranked movie at 22%. The top six movies cover 89.6% of the revenue, and the seventh-ranked movie collects less than 4% of tickets, with percentages declining thereafter.

⁹Lee (2006) examines the U.S. movies shown in Hong Kong and finds that the movies with a higher U.S. box office and action movies achieve a higher box office in Hong Kong. Kwak & Zhang (2011) report that, among the foreign movies shown in China, action and comedy movies enjoy a higher box office than drama movies.

Table 1: Movie Characteristics

Variables	Unweighted			Admission-Weighted		
	(1) All	(2) Domestic	(3) Foreign	(4) All	(5) Domestic	(6) Foreign
Age (Week)	7.06	7.57	5.66	7.71	9.73	5.33
RunTime (Minute)	101.9	98.45	111.5	117.2	110.7	125.0
Indicator variables:						
IMAX	0.10	0.03	0.29	0.42	0.18	0.70
3D	0.20	0.12	0.44	0.49	0.30	0.71
Foreign	0.27	0	1	0.46	0	1
Action	0.28	0.19	0.53	0.49	0.31	0.70
Comedy	0.31	0.35	0.21	0.26	0.35	0.16
Drama	0.33	0.35	0.28	0.34	0.47	0.19

Number of observations: 939, Foreign movies: 250, Domestic: 689.

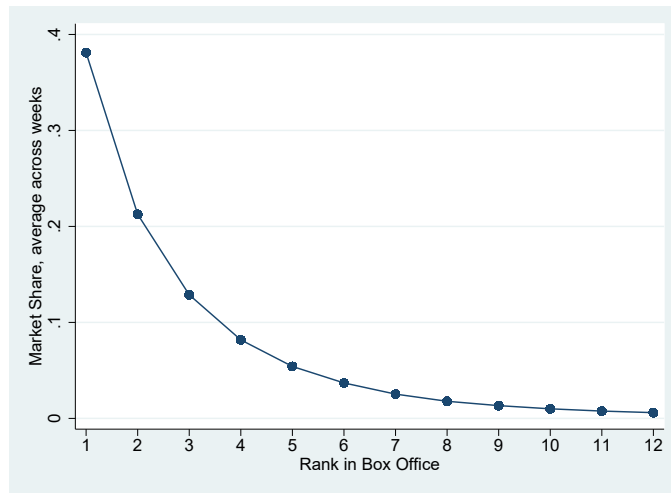


Figure 2: Average Share of Ticket Sales by Weekly Sales Rank




Figure 3: Age Profile of Movie Revenue

A common feature of box office revenue data is the steep drop-off in revenue that takes place from week to week. That is the case in our data as well. In order to see this, we perform a regression of the log of sales by movie and week on movie, year, month-of-year, and age fixed effects. For this regression, we use only the top six movies by box office revenue in each week, which enables us to drop movies that are re-released long after their original release. Age is defined as the number of weeks since the release of the movie, so there is a separate fixed effect for each age, up to eleven weeks (given our focus on the top six, there are only two movies in the data that make it to eleven weeks). Based on this regression, we predict sales for the average movie by week. For this prediction, we set the date to April 2012, make the prediction for every movie, and take the mean.¹⁰ The result appears in Figure 3. Predicted sales start around 95 million in the first week and drop to less than 50 million by week 3, and are under 10 million by week 5, with continued declines afterward.¹¹

4.1 Time-varying variables

Einav (2007) reports that there is seasonality in movie demand. Thus, we use a dummy variable for whether the current week has a holiday ($Holiday_t$) to capture the demand fluctuations of

¹⁰In order to account for the non-linear transformation in using a log regression to predict the level of sales, we use Duan’s smearing estimate. We use `levpredict` in Stata. See also Duan (1983).

¹¹Note that with product fixed effects, age and calendar date are not non-parametrically identified, so the fact that we restrict calendar date to enter by year and month-of-year is potentially important. We do not further explore the issue here, but it might be possible to exploit plausibly exogenous variation in release delays of foreign movies in China, similar to the way Mehta, Rysman & Simcoe (2010) use patent office delay in the context of patent citation age profiles.

movies within a month. The holidays included are New Year’s Day, Chinese New Year, Qingming Festival, May Day, Dragon Boat Festival, Mid-Autumn Day and National Day. On average, 20% of observations belong to movies showing on holiday.

We further include a linear time trend in the month to capture the dramatic increase in the Chinese movie market documented in Figure 1, as well as year fixed effects. The time trend captures issues such as growing income in China as well as growth in the number of theaters. We include a separate set of month-of-the-year dummies for foreign and domestic movies. Having two separate sets of month dummies is meant to capture anecdotal evidence that SARFT’s treatment of foreign movies varies by season.

4.2 Market Size and Market Share

In this subsection, we motivate several important modeling assumptions. An important restriction that we make for computational reasons is that consumers can select among six *named* movies in each week, in addition to a generic foreign and generic domestic outside option. We assume consumers can select among the six movies with the highest market share in each week.

There appears to be little gain to adding more named movies to the choice set. We calculate the box office share of the top six movies each week, the remaining foreign and the remaining domestic movies, and take the average over weeks. The results appear in Table 2. This table has 1,464 observations, which consists of the six top movies in each week and the two generic options (one foreign and one domestic) for 183 weeks.¹² We find that the top six movies have an average of 89.6% of the market. Thus, similar to what we saw in Figure 2, considering only the top six still captures most of the market. The generic foreign option gets about 3% and the generic domestic option gets about 7%. If we increased the top six movies to be the top ten, we capture 96.2%, an increase of less than 7 percentage points. Thus, there is little gain to expanding this number, and the computational cost would be high. Note that the characteristics presented in Table 1 are similar when using only the 427 movies that appear in the top six. For completeness, we recalculate Table 1 for these movies and present the results in Table 9 in Appendix A.

¹²For 16 weeks, we observe zero ticket sales for the foreign generic option, and we assume there was 1 ticket sold.

Table 2: Market Shares

Variables	(1) Mean	(2) SD	(3) Min	(4) Max
Box Office Share (%)				
Average movie	12.5	12.7	0.1	88.5
Top six movies	89.6	6.1	68.3	98.8
Other domestic movies	7.0	4.7	0.3	29.9
Other foreign movies	2.9	2.8	0.1	14.7
Market Share (% , out of potential market)				
Average over top six	0.5	0.6	0.0	8.5

1,448 observations

A potentially restrictive assumption in our model is that agents are myopic. While we consider a dynamic model for robustness, we believe the myopic model is reasonable. An important way in which agents might act dynamically is that they know when movies exit the theaters and make sure to see movies before that happens. However, in our data, for movies that are ever in the top six, the average percentage of their time that is spent in the top six is only 55.2%. That is, at the end of their time in the top six, movies do not disappear. Instead, they enter one of our generic options. Thus, consumers do not have to perceive movies in a dynamic way in order to be sure to see a given movie. In contrast, a movie's time in the top six accounts for most of its revenue: For movies ever in the top six, 85.7% of revenue is realized while in the top six. Weighted by ticket sales to emphasize top sellers, the average percentage of time spent in the top six is still only 69%, whereas the percentage of revenue realized while in the top six is 95.4%. Overall, we find these descriptive statistics consistent with our assumptions that consumers are not forward-looking, and that they choose among six top movies and two generic options (and the outside option).

In order to define market shares, we must define the potential market. We define China as a whole as the geographic market, which is analogous to Einav (2007) in the United States. Because movie theatres are often located in urban areas, we employ the population in an urban area instead of total population to measure the market size. We use the annual figure of total urban population in the year 2011, i.e. 354.256 million people, to measure the market size, and this size is denoted M . The population data is obtained from the China Statistical Yearbook. To compute market shares, we divide the ticket sales of movie j in week t by the market size. Let q_{jt} be the ticket sales

(quantity, not revenue) of movie j in week t . Then, $s_{jt}^{\text{data}} = q_{jt}/M$ is the market share of movie j . The outside good is defined as not watching a movie in a theater. The average market share of a movie is 0.5%, whereas the outside option averages 96%.

4.3 Data for Micro-moments

We employ summary statistics reported from a survey conducted by a Chinese consulting firm called Entgroup. The survey was conducted in February and March of 2013. The 6,027 respondents are consumers who had watched at least one movie in the theater in the previous year. The survey shows that 23.2% of the respondents watched 1-3 movies, 19.2% watched 4-6 movies, and 57.6% watched more than six movies in the previous year.

5 Model

This section presents our model for consumer demand for movies. It is meant to capture what we consider to be the three features of performance goods: rapidly exogenous evolution in choice sets, limited time to consume performances, and consumption durability. The limited time that consumers may allocate to performances is captured by assuming consumers can see at most one movie per week. Obviously, this is not strictly true, but we believe that it is a good representation of consumer decision-making. Consumption durability is captured by assuming that consumers see a given movie no more than once. We discuss relaxations of this assumption below.

In addition, we assume for now that consumers make their current choice myopically. This assumption might be problematic in some performance markets, but we believe it is reasonable in our setting. We discuss this assumption further below.

5.1 An overview

We present a simplified version of how the model works in Figure 4. The figure represents four time periods (weeks). The top row reports the time period and the set of exogenously available movies. For this example, we assume only two movies are available rather than six. In the first three periods, movies A and B are available. In the fourth period, movie A drops out, and movie C arrives. A consumer starts in period 1 having not seen any movies, and so starts with the choice




Figure 4: Simplified Representation of the Demand Model

set $\{A, B\}$. The three arrows from $\{A, B\}$ represent the three choices the consumer may make: the consumer can choose to see A , B or choose not to see a movie.

The exogenously available movies stay the same in period 2, so consumers will face one of three choice sets in period 2 depending on what they choose in period 1. Consumers that saw A are in the set $\{B\}$ in period 2, consumers that saw B are in the set $\{A\}$, and consumers that did not see a movie are again in $\{A, B\}$. Consumers can reach one for four states in period 3, because consumers that saw movies in both periods are now in state $\{\phi\}$, the empty set. These consumers cannot see a movie in period 3. In period 4, A drops out and C enters, so there are only two possible choice sets that consumers may reach in period 4: choice set $\{C\}$ for consumers that have already seen movie B in periods 1, 2, or 3, and choice set $\{B, C\}$ for consumers that have not yet seen B .

Figure 4 illustrates several points about our model. The set of potential choice sets evolves over time as movies exogenously enter and exit the market. If we think of the consumer's choice set as the consumer's state in a dynamic model, the number of states can grow from one period to the next, especially if there is no change in the available movies. However, turnover in the available movies typically leads to reductions in the number of potential states, and thus simplifies our computational

problem. Also, there are typically multiple paths by which a consumer may reach any given choice set. For instance, there are four arrows pointing to set $\{C\}$ in period 4, and there are multiple ways to reach each of the states that can lead to $\{C\}$.

In estimation, we assume a population of consumers starts in the first choice set in period 1, and then follows choice probabilities across each option (each arrow in Figure 4). Thus, we compute the share of the population that lands in each state in each period. Note that there is no simulation in this process. We compute the shares of consumers in each state exactly following the choice probabilities. In practice, we compute this for six movies per period rather than two, for three non-dynamic options (the two generic movies plus the outside option) rather than one, and for 183 time periods rather than four, so the problem is numerically challenging. In addition, we allow for persistent consumer heterogeneity in the form of permanent random coefficients, and this computation must be done separately for each consumer type. As described below, and as is standard, we use simulation to handle consumer heterogeneity.

5.2 The consumer problem

Now we present the model more formally. A continuum of consumers of size 1 indexed by i face discrete finite time. The set of all movies ever available can be indexed by j from 1 to J . In our case, $J = 427$. A subset of six of these movies is available in any given period. Denote the set of movies available in t as \mathbf{C}_t . We assume that \mathbf{C}_t follows an exogenous process. The set of six movies in \mathbf{C}_t can be combined into different choice sets. Denote the set of choice sets that can be reached by consumers as \mathcal{C}_t . The set \mathcal{C}_t has G_t elements, so G_t may be as high as 2^6 . We denote the elements of \mathcal{C}_t as C_{gt} , $g = 1, \dots, G_t$. In Figure 4, \mathbf{C}_t is the top row of a column, \mathcal{C}_t is a column, C_{gt} is each element of the column, and G_t is the number of elements in the column.

Denote the history of all movies seen by i up to period t as H_{it} . Let the function $C(H_{it}, \mathbf{C}_t)$ return consumer i 's choice set in t :

$$C(H_{it}, \mathbf{C}_t) = \{j : j \in \mathbf{C}_t, j \notin H_{it}\} \cup \{0, J+1, J+2\}.$$

The first part of the right-hand side says that consumers may choose among movies available in

the current period (that is, in \mathbf{C}_t) but that they have not seen previously (that is, not in H_{it}). The second part says that consumers always have three additional options. They may choose the outside option $j = 0$, or they may choose to see a generic foreign movie ($j = J + 1$) or a generic domestic movie ($j = J + 2$). These last options differ from the elements in \mathbf{C}_t in that they are always available and consumers may choose them repeatedly over time. Below, we also apply a simplified specification for utility for the generic movies. It must be that $C(H_{it}, \mathbf{C}_t) \in \mathcal{C}_t$, so $C(H_{it}, \mathbf{C}_t)$ must be equal to an element C_{gt} .

Let the utility to consumer i from choosing movie j in period t be denoted by u_{ijt} . Consumer i solves:

$$\max_{j \in C_{gt}} u_{ijt} \quad C_{gt} = C(H_{it}, \mathbf{C}_t).$$

We assume that utility takes on the functional form:

$$u_{ijt} = x_{jt}\beta + \xi_{jt} + \mu_{ijt} + \varepsilon_{ijt}.$$

The variables x_{jt} are K characteristics, observable to both the agent and the researcher. The characteristics reflect both movie characteristics, such as whether a movie is foreign, and calendar characteristics such as the month of the year and whether it is a holiday weekend. The scalar ξ_{jt} is observed by the agent but not the researcher. It represents unobserved quality, and will play the role of the econometric error term in our model. The term ε_{ijt} is distributed according to the Extreme Value distribution, and generates the familiar logit probability of choice. The term μ_{ijt} represents the consumer match to the product based on observable characteristics. Following Berry (1994) and Berry et al. (1995), we specify it as:

$$\mu_{ijt} = \sum_{k=1}^K x_{jkt} \sigma_k \nu_{ik}$$

where $\nu_{ik} \sim \mathcal{N}(0, 1)$. Thus, ν_{ik} captures consumer heterogeneity over preferences for observable characteristics such as whether a movie is foreign and whether it is enhanced with features such as IMAX filming. The parameters β and σ_k , $k = 1, \dots, K$ are to be estimated. We refer to them together as $\theta = \{\beta, \{\sigma_k\}_{k=1, \dots, K}\}$. Furthermore, for convenience, we denote the mean utility of product j in period t as $\delta_{jt} = x_{jt}\beta + \xi_{jt}$.

5.3 Market shares

Given these assumptions, the conditional probability $P_{ijt}(C_{gt})$, the probability of i choosing j in t conditional on having choice set C_{gt} , is:

$$P_{ijt}(C_{gt}) = \begin{cases} \frac{\exp(\delta_{jt} + \mu_{ijt})}{\sum_{k \in C_{gt}} \exp(\delta_{kt} + \mu_{ikt})} & j \in C_{gt} \\ 0 & \text{otherwise} \end{cases}. \quad (1)$$

In Figure 4, $P_{ijt}(C_{gt})$ is the probability of being on each arrow leading from a given choice set.

As is clear from Figure 4, there may be multiple choices that lead from one choice set to another. Let $B_{gg't}$ be the set of products j such that choosing one leads from choice set g in period t to choice set g' in period $t+1$. The set $B_{gg't}$ accounts for the deletion of j from C_{gt} , and any products that enter or exit \mathbf{C}_t :

$$B_{gg't} = \left\{ j : C_{g't+1} = \underbrace{(C_{gt} \setminus \{j\})}_{\text{current}} \cup \underbrace{(\mathbf{C}_{t+1} \setminus \mathbf{C}_t)}_{\text{entering}} \setminus \underbrace{(\mathbf{C}_t \setminus \mathbf{C}_{t+1})}_{\text{exiting}} \right\}. \quad (2)$$

Let s_{igt} be the share of consumers of type i with choice set g in period t . Thus, $\sum_{g=1}^{G_t} s_{igt} = 1$. We refer to s_{igt} as the unconditional probability or unconditional share. To compute s_{igt} , we assume that there is only one possible choice set in the first period: $\mathbf{C}_1 = \{\mathbf{C}_1\}$. Thus, $s_{i11} = 1$ for all i . Unconditional shares evolve as follows:

$$s_{ig't+1} = \sum_{g=1}^{G_t} \sum_{j \in B_{gg't}} P_{ijt}(C_{gt}) s_{igt} \quad \forall g' = 1, \dots, G_{t+1}, t = 1, \dots, T-1. \quad (3)$$

In the data, we observe the unconditional share of consumers choosing each product j in each period t . Our model defines that as:

$$\hat{s}_{jt} = \int \sum_{g=1}^{G_t} P_{ijt}(C_{gt}) s_{igt} f(i) di \quad (4)$$

where $f(i)$ is the distribution of consumer types i , assumed to be the multivariate normal distribution.

5.4 Forward-looking behavior

In some performance goods settings, our assumption of myopic behavior may not be reasonable.¹³ In this sub-section, we provide a model that allows for forward-looking behavior. We assume consumers have perfect foresight over all future values of δ_{jt} but not over ε_{ijt} . That is, consumers know all the movies that will arrive and leave, and the mean utilities that the movies will provide. Perfect foresight is a strong assumption, but we believe that to the extent that forward-looking behavior might be important, it is because consumers know that particular movies are arriving or leaving.

The inclusive value represents the value that a consumer expects when they face a given choice set. Under our logit assumptions, the inclusive value has a convenient closed-form. Define the inclusive value from making a choice from set g in period t to be:

$$V_{igt} = \ln \left(\sum_{j \in C_{gt}} \exp(\delta_{jt} + \mu_{ijt} + \lambda V_{ig't+1}) \right).$$

where g' is the choice set in $t+1$ that a consumer will realize when they start in g, t and pick j (which is written out formally in the brackets in Equation 2). The variable λ is the discount rate. For this calculation, we assume that $V_{igT+1} = 0$ for all i and g . Thus, for a consumer in the final period T , the choice problem is the same whether we use myopic or forward-looking behavior.

Thus, we can define the utility to i from movie j as:

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \lambda V_{ig't+1} + \varepsilon_{ijt}$$

Rewriting Equation 1, the new choice probability is:

$$P_{ijt}(C_{gt}) = \begin{cases} \frac{\exp(\delta_{jt} + \mu_{ijt} + \lambda V_{ig't})}{\sum_{k \in C_{gt}} \exp(\delta_{kt} + \mu_{ikt} + \lambda V_{i\tilde{g}'t})} & j \in C_{gt} \\ 0 & \text{otherwise} \end{cases}.$$

Here, we write \tilde{g}' in the denominator to distinguish it from g' in the numerator, as different choices will lead the consumer to different choice sets. The rest of the model, such as the determination of s_{igt} , remains the same. We can estimate this model by backward induction. For a given guess of the

¹³For example, we understand from private conversations with staff at the Museum of Fine Arts in Boston that when the museum announces that a temporary exhibit will be closing, attendance at that exhibit increases. That is evidence of dynamic behavior in exhibit attendance.

parameters, we can calculate the utility and probability of each choice in the last period. We can then calculate the utility and probabilities of each choice in period $T - 1$, accounting for the continuation value associated with each choice. We proceed backward through each period sequentially.

We wish to estimate the discount rate. Magnac & Thesmar (2002) argues that identification of the discount rate requires variation in the continuation value that is not reflected in the current values. Perfect foresight models generate this kind of variation naturally. In our setting, movies that arrive or leave in future periods affect the future payoff but not otherwise the current payoffs.

In considering the discount rate, it is important to recognize that the discount rate we estimate is unlikely to correspond to the time value of money. The discount rate in our model reflects how consumers adjust movie-going this week to changes in movie availability next week. Our prior belief is that consumers heavily discount this continuation value, and indeed we find it to be so below.

We focus on the perfect foresight model not only because we believe that it well-captures the issues that concern us, but also because it is computationally straightforward to estimate. The perfect foresight model requires no further assumptions and does not require a fixed-point algorithm. In contrast, limited information models typically require an assumption of stationarity as well as assumptions on the information set that consumers have. Researchers may wish to invoke Inclusive Value Sufficiency (as in Gowrisankaran & Rysman, 2012) but that introduces multiple fixed point algorithms, as well as questions about how to discretize or otherwise approximate the state space.

5.5 Multiple purchase

We briefly describe how we would extend the model to relax the assumption that consumers see a movie only once. We do not estimate this model, but the extension is useful in order to understand the model. It would be relatively easy to allow consumers to see a movie multiple times with decreased utility. Intuitively, thinking of Figure 4, consider a consumer in period 2 who has already seen movie A . The consumer is in set $\{B\}$, where A is not allowed. But the important feature of set $\{B\}$ is not that A is not allowed, but rather that the consumer has already seen A . It would be straightforward to allow a consumer with choice set $\{B\}$ to choose between both A and B , but assign A some reduction in utility, presumably a parameter to be estimated. We could assume

that choosing to see a movie multiple times does not further reduce the utility of the movie after the first viewing. As a result, this extension does not affect the overall dynamic process in our model, and thus would not be more difficult to estimate. This approach would be appealing if we had data on how often consumers saw individual movies, and the data showed that multiple viewings were important.

Note that this approach would assign the same utility reduction to each viewing of a movie after the first one. That is, the consumer would get the same utility from seeing a movie the second, third, and fourth time. In some settings, it might be more natural to assume that consumers experienced further declines in utility the more times the consumer saw a performance. That would be a more significant extension to our model in terms of computational difficulty, but we believe our model provides a good template for how to approach this problem.

6 Estimation

This section discusses estimation of the model. We first discuss the approach to the aggregate sales data, and then discuss micro-moments.

6.1 Aggregate data moments

First, we cannot compute Equation 4 analytically. For this step, we use simulation. We draw S values of ν_{ik}^s , $k = 1, \dots, K$ and $s = 1, \dots, S$.¹⁴

For a given set of parameters θ and a guess of mean utilities δ_{jt} , we compute $P_{ijt}(C_{gt})$ and then s_{igt} for each C_{gt} and movie j in the model, as described above in Equations 1 and 3. We do so separately for each draw of ν^s . We then replace Equation 4 with the discrete equivalent. In order to emphasize the dependence of the predicted market share on parameters and mean utilities, we write $\widehat{s}_{jt}(\theta, \delta)$, where δ is the vector of elements δ_{jt} .

As in Berry et al. (1995), we recover δ for any set of parameters θ via the fixed point equation:

¹⁴Using s for both samples and market shares is somewhat confusing, but it is clear in context. In practice, we try two sampling schemes. In the first, we draw ν_{ik}^s from a Halton sequence, setting $S = 300$. A Halton sequence produces an even spread of draws across percentiles of the normal distribution. In the second, we use importance sampling to overweight draws of ν_{ik}^s that are likely to attend movies, again using 300 draws. Details are available upon request. Results are similar, as we show below.

$$\delta'_{jt} = \delta_{jt} + \ln(s_{jt}^{\text{data}}) - \ln(\widehat{s}_{jt}(\theta, \boldsymbol{\delta})).$$

As above, s_{jt}^{data} are the market shares observed in the data. For any guess of parameters θ , we solve this equation by successive approximation. That is, we plug in a guess of $\boldsymbol{\delta}$, compute $\boldsymbol{\delta}'$ and iterate until convergence. Note that the theorem in Berry (1994) that the fixed point equation is a contraction mapping does not necessarily apply to dynamic models. As in Gowrisankaran & Rysman (2012), our method is appropriate only under the assumption that the solution is unique. We have not experienced any problems with multiple solutions in practice.

Based on the solution to the fixed point equation, we compute the econometric error term as:

$$\xi_{jt} = \delta_{jt} - x_{jt}\beta$$

and we assume a set of instrumental variables Z_{jt} is exogenous such that $E[m_1(\theta)] = E[\xi_{jt}|Z_{jt}] = 0$. We estimate via two-step GMM. We always assume that in the first period of the data, no consumers have seen a movie. In order to address this initial conditions problem, we drop the first four weeks of data in forming our moments. As there is frequent turnover in which movies are available, the “burn-in” period before consumers are reasonably distributed across choice sets is relatively short. As a result of dropping these observations, the tables of results in Section 7 list 1,432 observations rather than 1,464.¹⁵

In practice, we include a full set of movie fixed effects, so we do not estimate β for any variables that do not vary over time. The generic domestic and foreign outside option each have a dummy variable indicating their type, and are subject to the time-varying explanatory variables (time trend, holiday, and month-of-year effects) but are not further parameterized. Our base specification places random coefficients on three variables: the constant term, a dummy for whether a movie is foreign, and a dummy for whether a movie filmed in either IMAX or 3D, which we call *enhanced*. These are the most important variables for our research question. We experiment with other specifications as well.

¹⁵A potential problem for our method is if movies appear in the top six in non-adjacent weeks, because as written, we would lose track of which consumers have seen the movie. That happens only once in our data set. One movie leaves the top six in week 75 and returns in week 76. We keep this movie in \mathcal{C}_t for $t = 75$, and drop the sixth most popular movie from \mathcal{C}_{75} .

We assume that all explanatory variables are exogenous. Recall that price is not an explanatory variable. However, the presence of consumer heterogeneity terms (σ_k) means we still need additional instrumental variables to achieve identification. Our first set of instruments follows Berry et al. (1995). Because we take product introductions as exogenous, we use sums over the characteristics of other movies in the top six in the same week. For this calculation, we use the following variables: dummies for whether the movie is enhanced (3D or IMAX), foreign, action, comedy or drama, and the number of weeks since the movie’s Chinese release, the number of weeks since the movie’s international release (set to 0 for Chinese movies), and the movie’s runtime. Additionally, Gandhi & Houde (2019) recommend instruments that emphasize how differentiated a product is from others on the market. We construct these for the instruments based on dummy variables. We do so by interacting the Berry et al. (1995) instruments with the dummy variable in question. Thus, the sum over the *enhanced* dummy will be interacted with whether the movie in question is enhanced, and so will be high only for enhanced movies.

6.2 Incorporating the Micro-moments

To improve the estimation, we incorporate two micro-moment conditions based on the survey data. Specifically, we use the information that, conditional on watching at least one movie, the probability of watching 1-3 movies is 23.5% and the probability of watching 4-6 movies is 19.2%.

In order to compute the predictions of these variables from our model, we augment the state space for consumers to track not only which movies they have seen, but also how many times they have been to the movies. That is, we denote the state of a consumer as $\{C_{gt}, n_{it}\}$ where n_{it} is the number of movies that i has seen in the previous year. When a consumer chooses to see a movie and $n_{it} < 7$, then $n_{it+1} = n_{it} + 1$. Intuitively, we duplicate Figure 4 seven times, and as the population of consumers moves across the figure, the ones that see movies also move from figure to figure. We track this only for the 12 month period leading up to the observation of our moment (January 2013), not for the entire 183 week period of the data. We assume that n_{it} takes on a maximum of 7 to reflect our survey data, although a consumer with $n_{it} = 7$ can continue to go to the movies.

To be clear, this new state variable does not affect consumer decision-making. The consumer still

cares only about C_{gt} . Tracking n_{it} allows us to form predictions that may be compared to the survey data. In particular, at $t = 57$, we compute P_{in} , the probability that consumer i saw n movies in the previous year, for $n = 0, 1, \dots, 6, 7+$. Conditional on watching at least one movie, the probability of watching 1-3 movies is then $P_{i1-3} = \sum_{n=1}^3 P_{in} / \sum_{n=1}^{7+} P_{in}$ and the probability of seeing 4-6 movies is $P_{i4-6} = \sum_{n=4}^6 P_{in} / \sum_{n=1}^{7+} P_{in}$. We take the average of P_{i1-3} and P_{i4-6} , i.e. $P_{1-3} = \frac{1}{S} \sum_{i=1}^S P_{i1-3}$ and $P_{4-6} = \frac{1}{S} \sum_{i=1}^S P_{i4-6}$. We postulate the micro-moment conditions as follows

$$E[m_2(\theta)] = E \left[\begin{array}{c} P_{1-3}^{\text{data}} - P_{1-3}(\theta) \\ P_{4-6}^{\text{data}} - P_{4-6}(\theta) \end{array} \right] = 0, \quad (5)$$

The variables on the left are the probabilities observed in the survey data. Thus, the stacked moment conditions are

$$E[m(\theta)] = E \left[\begin{array}{c} m_1(\theta) \\ m_2(\theta) \end{array} \right] = 0. \quad (6)$$

Here, $m_1(\theta)$ are the aggregate-data moments, as discussed in Section 6.1. The GMM estimator given our stacked moment conditions is defined as $\min_{\theta} E[m(\theta)]' \Omega E[m(\theta)]$. We follow the two-step procedure for GMM estimation proposed in Hansen (1982) and initialize it with an identity matrix as the weighting matrix Ω . We draw a new sample of draws ν_k^s for the micro-moment calculation. Thus, the weighting matrix is block-diagonal as in Petrin (2002). In the second stage of the GMM optimization routine, the weighting matrix of the micro-moment conditions is computed using a variance-covariance matrix of the micro-moment conditions.

It is difficult to know how to weight the two sets of moments in estimation. Although formally, the survey data has more observations, we believe it is less reliable than the administrative ticket data. Following Li, Mazur, Park, Roberts, Sweeting & Zhang (2019), we impose that the two sets of moments are weighted equally. Formally, we impose that the sum of the weights within each set of moments (the aggregate-data moments and the micro-moments) are equal. In the second stage of GMM, we impose that the weighting matrix is diagonal and we allow the relative weights within each set of moments to reflect the relative inverse of the variance of the moment, but we still normalize so that the weight on each set of moments is equal.

We finish this section with a brief heuristic discussion of identification. Our model has the same parameters as Berry et al. (1995), and thus we can think of identification in a similar way.

Products with different characteristics attract different levels of market share. To the extent the product draws market share from products with a similar characteristic, estimation will find that the random coefficient on that characteristic has a large variance. To the extent that the product draws market share proportionally from all products, the variance will be small as in a standard logit model. Like Gowrisankaran & Rysman (2012) (and unlike Berry et al., 1995), the model makes use of market shares over time. In the period following a popular foreign movie, seeing high market share move to competing foreign films indicates that the random coefficient on *foreign* has high variance. That is, consumers that like foreign movies saw the first foreign movie and then, when it was out of the choice set, those consumers switched to another. It is difficult to design reduced-form models that capture this appropriately because the market share of a movie in any period depends on how long it has been on the market, and which movies it faced in those periods (and by extension, which movies those rivals faced in their time on the market, and so on). An advantage of a structural model in a setting like this is that it resolves these complex interactions in a coherent and parsimonious way.

7 Empirical Results

This section discusses the empirical results obtained from the demand model described in the previous section. Column 1 of Table 3 reports estimates from using a standard Berry et al. (1995) model, i.e. a static random coefficients model. A striking feature of column 1 is the large negative and significant age trend. Static models can match the kind of declines in market share that we see in the data (as evidenced in Figure 3) only with a strong reduced-form age profile. Column 2 adds the micro-moment but the age trend is almost unchanged. The static model has no mechanism for matching the micro-moment and so it does not qualitatively affect the results (although the coefficient on *foreign* drops substantially).

Column 3 estimates our dynamic model, in particular the model with myopic consumers who experience consumption durability. The coefficient on *age* drops almost in half. Our preferred specification is Column 4, in which the micro-moments are imposed on the dynamic model. This

Table 3: Demand Estimates

Parameters		Static	Static	Dynamic	Dynamic	Dynamic
Non-linear	Constant	1.645 (0.132)	0.994 (0.015)	2.427 (0.122)	10.526 (0.138)	10.632 (0.081)
	Enhanced(3D or IMAX)	0.881 (0.169)	0.353 (0.165)	0.038 (0.044)	1.531 (0.097)	1.604 (0.071)
	Foreign	3.144 (0.072)	0.013 (0.019)	4.678 (0.078)	3.343 (0.100)	3.365 (0.075)
Linear	Age	-0.607 (0.056)	-0.545 (0.062)	-0.312 (0.069)	0.006 (0.076)	-0.080 (0.063)
	Holiday	0.405 (0.388)	0.404 (0.415)	0.450 (0.387)	1.300 (0.401)	1.435 (0.397)
	Consumption Durability	No	No	Yes	Yes	Yes
	Micro-moments	No	Yes	No	Yes	Yes
	Named movies	6	6	6	6	7
	Observations	1,464	1,464	1,464	1,464	1,647

Specifications include movie fixed effects, month-of-year fixed effects separately for foreign and domestic movies, year fixed effects and a month time trend. *Named movies* refers to the number of movies in a week that consumers track whether they have seen. Increasing to seven in the last column increases the number of observations.

specification leads to a dramatic increase in the random coefficient on the constant term. That is, the way to match the repeat viewing in the survey data is to greatly increase consumer heterogeneity, so some consumers highly value going to the movies and go repeatedly. Because the effective movie-going population is much smaller in this specification, the age profile is no longer necessary to create the drop-off in sales with age. The coefficient on *age* is insignificant and close to zero in magnitude. Thus, despite the enormous age effects in the raw data (as evidenced in Figure 3), the age profile can be entirely explained by the consumption durability of movie consumption.¹⁶

In thinking about identification, note that consumption durability does not necessarily imply that the *age* coefficient would be zero. Consumption durability implies demand will fall with time, but observed sales could be higher or lower than consumption durability would predict. For instance, the *age* coefficient in column 3 is negative. However, when the micro-moments are imposed, we see that coefficient on *age* driven to zero.

In Column 4, the random coefficient parameters on movies *enhanced* and *foreign* are also statistically significant. The parameter on *foreign* in particular is fairly large. That will drive our result

¹⁶The age trend is not separately identified from movie fixed effects and a week time trend, which is one reason we use a month trend. The results are robust to alternative treatments of the calendar time effects, such as using only year dummies.

in the next section that there is relatively muted substitution between foreign and domestic movies.

To establish robustness, we consider several alternative models. It is natural to wonder whether our assumption that consumers track only six movies per period is restrictive. To consider this, we estimate the model allowing consumers to track seven movies. Results appear in column 5 of Table 3. They appear very similar to the results in column 4 that uses six movies. That is not surprising given the low market shares associated with low ranked movies.

Another possible concern is that, as described in Section 6, we construct our moments based on the assumption that $E[\xi_{jt}|Z_{jt}] = 0$. However, it might be more natural in a dynamic framework to assume that $E[\xi_{jt} - \xi_{jt-1}|Z_{jt}] = 0$. This is the approach of Lee (2013). This “differenced” model focuses on changes over time rather than levels. For this specification, we also first-difference the instruments. The effect of changing from levels to differences is muted in our case because we have product fixed effects in the levels model. It is analogous to switching from fixed effects to first differences in a linear panel data estimation setting, which are asymptotically identical when using a flexible weighting matrix, such as we do. Not surprisingly, we find similar results. These appear in the Table 10 in Appendix A.

We also consider a model in which consumers have perfect foresight as to what movies will be available, as described in Section 5.4. We perform a grid search over values of the discount rate λ from 0 to 1. The specification is otherwise identical to that in Column 4 of Table 3, which includes the micro-moment. For each value of λ , we estimate the rest of the parameters as above. We find that the objective function is minimized at a discount rate of $\lambda = 0$.¹⁷ Intuitively, consumers do not respond this week to the future availability of movies, which we believe is a reasonable result. Thus, our assumption of myopia fits the data well. Parameters are naturally the same as in column 4 but standard errors change because of the extra parameter.¹⁸ This result appears in Table 10 in Appendix A.

Interestingly, for higher values of the discount rate (that is, more utility weight on the continuation value) we find a more negative age profile. For instance, when the discount rate λ is set to

¹⁷Around $\lambda = 0$, we consider increments in the grid search as low as 5×10^{-5} .

¹⁸We find a standard deviation of λ of 0.023. We calculate the standard error with the usual sandwich estimator for optimal GMM. In this calculation, we do not address the issue that the parameter is on an inequality constraint.

0.5, we find the age profile coefficient increases in magnitude to -0.19 and is statistically significant. Intuitively, the model finds that forward-looking consumers see movies earlier because consumers anticipate that a movie will decline in value. Whereas consumption durability reduces the importance of the age profile, forward-looking behavior can increase its importance. Given our estimate of λ , this point does not affect our evaluation of this market, but it may be interesting for other work on performance goods. The result appears in the Appendix (Appendix A, Table 10),

As an additional robustness check on this issue, we estimate a model with no discounting (i.e. $\lambda = 1$) but in which consumers look only one period into the future. This might be a realistic approximation of forward-looking behavior in the market for movies. In Table 10, results appear similar to the case with perfect foresight. One further issue we consider is our scheme for integrating over consumer heterogeneity. Our main results use Halton sequences. When we use importance sampling, we find similar results, as shown in Table 10 (see also Footnote 14).

Next, we regress the movie-specific effects from the demand estimation on time-invariant movie characteristics and report the results in Table 4. Focusing on column 4, our preferred specification, we see that enhanced movies and action movies have positive and significant coefficients. We further control for the weeks since international release, and it is negative. Thus, Chinese consumers are more likely to see movies released close to their international release. This may be because there is significant marketing close to the release day, or because release delay allows counterfeit versions of the movie to reach consumers. The coefficient on being foreign is insignificant, but the coefficient is difficult to interpret because of the separate foreign and domestic month-of-year fixed effects in the demand specification. The choice of which month to exclude from the month-of-year fixed effects in the estimation from Table 3 can greatly affect the coefficient on the indicator for being a foreign movie in Table 4.

8 Counterfactual Experiments

Since 2012, China has agreed to increase the import quota for foreign movies from 20 to 34 in each year. The import liberalization specifies an extra 14 foreign movies in 3D or IMAX formats,

Table 4: Regression of Movie Fixed Effects on Movie Characteristics

	Static	Static	Dynamic	Dyanamic
Enhanced	0.122 (0.277)	0.320 (0.247)	0.739 (0.184)	1.082 (0.220)
Foreign	-0.484 (0.300)	-0.180 (0.268)	-0.854 (0.199)	-0.298 (0.238)
Weeks since int'l release	-0.006 (0.004)	-0.005 (0.004)	-0.011 (0.003)	-0.013 (0.004)
Run time (log)	2.711 (0.814)	2.547 (0.726)	3.336 (0.540)	4.964 (0.646)
Action	0.192 (0.267)	0.166 (0.238)	0.303 (0.177)	0.407 (0.212)
Comedy	0.158 (0.284)	0.171 (0.253)	0.205 (0.188)	0.311 (0.225)
Drama	-0.087 (0.273)	-0.064 (0.244)	-0.079 (0.181)	0.074 (0.217)
Constant	-12.563 (3.759)	-12.019 (3.355)	-15.602 (2.495)	-23.777 (2.986)

418 observations. The columns are defined analogously to Table 3. The first two columns implement a static demand model as BLP. The next two columns add consumption durability. The second and fourth columns implement micro-moments.

which are mainly produced in the United States. This section performs counterfactual experiments to evaluate this import liberalization on consumer and producer welfare. An assumption we make to perform these counterfactual experiments is that the producers do not revise the attributes of their movies in response to the import liberalization.

We first discuss how to select the counterfactual set of 20 movies, and then present results.

8.1 A model of movie selection by SARFT

We consider several models of which 20 movies would have been selected if there had not been a liberalization from 34 to 20 movies. We compare taking the 20 movies from the bottom of the admissions distribution, from the top of the admissions distribution, and from an empirical model designed to estimate how the Chinese government chooses which movies to select. In this subsection, we discuss this model and its predictions.

Our approach to modeling SARFT's decision-making is to form a list of the top 100 movies per year by North American box office, and then perform a probit regression movie-by-movie on which ones are accepted into China under revenue-sharing. We then use simulation techniques to construct

an ordering of the movies by the latent value in the probit model, and assume that SARFT would select in order of this latent variable.

In more detail, we select the top 100 movies by North American box office revenue according to Boxoffice Mojo.com. We do so for each year from 2008-2015, which allows us to study SARFT decisions before and after the policy change in 2012. We estimate two separate regressions for before and after the policy change. We assign movies to years based on the date of release in North America and use their lifelong revenue, so even a movie released in late December may be among the highest earners. We assume the selection process for each movie is governed by a Probit model:

$$Y^* = z\alpha + \eta, \quad Y = \mathbb{1}\{Y^* \geq 0\}.$$

where $\eta \sim \mathcal{N}(0, 1)$ and $Y = 1$ if the movie is selected for revenue sharing by SARFT. In selecting the top 100 movies in a year, we do not include movies that enter by co-production. Also, we assume that movies that enter by fixed fee would have entered by revenue-sharing if they could have. Thus, we assign movies that enter by fixed fee to have $Y = 0$.

In considering what variables should be included in z , North American revenue appears to be of primary importance. However, there are other factors. To see this, consider Table 5, which reports the share of top movies (ordered by North American box office) selected for revenue-sharing by SARFT.¹⁹ Practically every movie comes from the top 100. The share coming from the top 50 though is always less than double the share coming from the top 100, so some movies are being selected from outside the top 50. While SARFT selected 7 or more of the top 10 in every year since 2010, SARFT selects all of the top 10 movies in only one year, 2015. That is, SARFT is regularly passing on movies in the top 10 and top 50.²⁰

For explanatory variables, we use the log of North American box office, runtime, genre indicators, and indicators for whether the movie is IMAX, 3D, rated R, and the calendar year.²¹ We also

¹⁹Note that movies are organized by release year, not selection year. So for instance, 25 of the top 100 movies are selected from the 2011 release year even though only 20 movies are selected in 2011 because some 2011 movies are selected in 2012, when the quota increased to 34.

²⁰For movies with Chinese release dates close to their U.S. release dates, North American revenue would be unknown the SARFT at the time of their decision. We ignore this issue. Box office revenue can often be predicted with at least some level of accuracy.

²¹We do not attempt to measure other factors that SARFT appears to account for, such as whether the movie glorifies foreign military, or is about religion or the occult. These are difficult to quantify. We briefly explored a

Table 5: Share of Top Movies That are Selected for Revenue-Sharing by SARFT

Year	% of Top 100	% of Top 50	% of Top 10
2008	19	32	50
2009	22	40	70
2010	21	36	70
2011	25	42	90
2012	30	52	80
2013	27	42	70
2014	29	50	80
2015	27	48	100

Top movies are ordered by North American box office revenue.

include dummy variables for the nationality of production. To do so, we partition movies into four categories. Movies can have 1) U.S.-only producers, 2) non-US, non-Chinese production, 3) joint U.S. and non-U.S. producers without Chinese producers or 4) joint foreign and Chinese producers. In our regressions, we include indicator variables for categories 2, 3, and 4, so the indicator for category 1 is the excluded variable. In the earlier time period, there are five movies in our data set in category 2 and all of them receive revenue-sharing. We drop these movies from our probit regression, so the earlier period has 395 rather than 400 observations. As we discuss below, we account for this feature in how we compute counterfactual outcomes.

Results from estimating the probit model appear in Table 6. We provide separate regressions for the four years after the policy, 2012-2015, and the four years before, 2008-2011. Starting with column 1, which gives results for the 2012-2015 period, we see that box office revenue, IMAX, 3D, and being an action movie are all strong positive predictors of selection, whereas being rated R is negative, all as expected. These parameters generate reasonable magnitudes of these effects. For instance, the effect of increasing box office by 1% is to increase the probability of selection by 0.2 percentage points.²²

machine-learning approach based on movie reviews, but with only a limited number for observations in each regression, and with much of the outcome explained by box office revenue, it did not appear that we had a data set suitable for such an approach.

²²We drop two movies, *Iron Man 3* and *Looper* that aimed for co-production status but were turned down very late in the process after they were produced, and entered China on a fixed-fee contract. Arguably, these movies never were considered for revenue-sharing and should not be included in this regression. When we include these two observations, results are almost identical. One change is that the coefficient on *Action* decreases slightly, as these are both action movies that are recorded as $Y = 0$.

Although the quota on foreign movies restricts consumer choice, these results suggest that conditional on the quota, SARFT selects movies that are popular with consumers. This may have been driven by concerns about consumer welfare, as well as revenue concerns, as SARFT collects a per ticket tax, and profits further through its ownership of CFG. SARFT may also value Chinese influence over international production decisions. SARFT's emphasis on popular movies suggests SARFT is not primarily attempting to protect domestic producers. However, as pointed out above, SARFT passes on some movies that seem likely to be popular. These decisions often have political explanations. Exactly how SARFT trades off these issues is interesting, but beyond the scope of this paper.

Given the primacy of box office revenue, we consider a model that adds the square of log box office revenue in column 2. Results for the other variables are quite similar. We find a concave effect for box office revenue, with a peak at about 90th percentile of box office revenue. The fit of the models is quite similar in terms of which models are selected (results available upon request).

As an alternative, we also estimate the Probit model using the period before the policy change, when only 20 movies were selected. Results are in column 3. Box office appears less important, with a coefficient about two-thirds of what we found in the post-period, and a marginal effect of only 0.075 percentage points. Other coefficients, such as on 3D, Action, and Rated R, are all about 50% larger. Indeed, although the liberalization specified that the increase from 20 to 34 movies was to include 14 movies that were 3D or IMAX, we do not find the constraint on 3D and IMAX to be binding. As we can see from column 3, there was a significant preference for IMAX and 3D movies before the policy change. Note that the lower coefficient on box office revenue in the pre-2012 period reflects the joint decisions of both SARFT and movie producers. Producers of internationally popular movies may have been less eager to release movies in China when its box office tended to generate less revenue and piracy was more of an issue.

Now we turn to selecting which movies would have been selected by SARFT if there had been no liberalization. We assume the 20 movies it would have selected come from the 34 movies it actually did select under liberalization. In both the pre and post period, SARFT stated that it would select

Table 6: Probit Model of SARFT's Decision to Select Movies for Revenue-Sharing

	Period 2012-15		Period 2008-2011	
ln(Box Office)	0.841	16.230	0.544	18.700
	(0.151)	(6.427)	(0.177)	(8.079)
ln(Box Office) ²		-0.418		-0.493
		(0.17)		(0.219)
U.S. and non-U.S. producer	0.183	0.204	0.503	0.520
	(0.197)	(0.199)	(0.219)	(0.225)
Chinese producer involved	0.906	1.014	0.651	0.543
	(0.537)	(0.555)	(0.945)	(0.942)
No U.S. producer	0.288	0.254	N/A	N/A
	(0.890)	(0.960)		
Rated R	-1.074	-1.134	-1.605	-1.703
	(0.247)	(0.253)	(0.386)	(0.405)
IMAX	0.509	0.554	0.533	0.607
	(0.237)	(0.241)	(0.261)	(0.263)
3D	0.868	0.937	1.347	1.376
	(0.215)	(0.220)	(0.308)	(0.315)
ln(RunTime)	1.222	1.244	2.927	3.077
	(0.694)	(0.700)	(0.854)	(0.875)
Action	0.593	0.613	0.990	0.999
	(0.23)	(0.24)	(0.23)	(0.238)
Comedy	-0.116	-0.124	-0.414	-0.448
	(0.25)	(0.25)	(0.25)	(0.251)
Drama	-0.233	-0.202	-0.711	-0.722
	(0.23)	(0.23)	(0.26)	(0.266)
2009			-0.003	0.009
			(0.28)	(0.290)
2010			-0.449	-0.473
			(0.31)	(0.313)
2011			-0.566	-0.657
			(0.31)	(0.320)
2013	-0.163	-0.142		
	(0.26)	(0.26)		
2014	-0.116	-0.156		
	(0.26)	(0.26)		
2015	-0.166	-0.139		
	(0.27)	(0.27)		
Constant	-21.970	-163.600	-24.540	-192.400
	(3.75)	(59.46)	(4.33)	(75.020)
Observations	398	398	395	395

The top 100 movies by the lifetime of North American box office revenue among all movies released in each of four years. The dependent variable is equal to 1 if the movie was selected for revenue sharing. Standard errors are in parenthesis. For 2012-2015, we drop two movies that were meant for co-production. For 2008-2011, we drop five movies with no U.S. production, as all of these are selected for revenue-sharing. For the nationality of production, *Only U.S.* is excluded. *U.S. and non-U.S. producer* indicates a non-U.S., non-Chinese producer working with a U.S. producer. *Chinese producer* indicates a Chinese producer working with a U.S. or non-U.S. producer. *No U.S. producer* indicates no U.S. producer and no Chinese producer.

at least six non-U.S. movies. In fact, we observe less than six movies without the involvement of a U.S. firm in all but one year, 2014. Perhaps SARFT counts movies with at least some non-U.S. production, of which there are many. We assume that all of the entirely non-U.S. produced movies that appear in the 34 for a given year are selected to be among the 20. From 2012 to 2015, this accounts for 4, 5, 6, and 2 movies.²³

Second, we construct $z\hat{\alpha}$ for each of the remaining 34 movies, where $\hat{\alpha}$ comes from the Probit estimation in column 1 of Table 6. We draw values of η from the standard normal distribution, and thus simulate an ordering of SARFT's preferences over the movies it can select from. We assume that it fills up what remains of the 20 slots based on this preference ordering. We repeat this process 100 times, thus generating a distribution of the set of 20 movies that SARFT would have selected without liberalization. In order to better understand our model and our results, we also present results if instead of using the ordering from the Probit model, we assumed that SARFT selected from its set of 34 strictly based on admissions. We consider both orderings, from the top and bottom of the admissions variable.

In constructing these sets, we ignore that the constraint to pick a certain number of IMAX and 3D movies existed only under liberalization. As we stated above, the constraint did not appear binding and SARFT exhibited a strong preference for such movies even before the constraint existed. Another potential drawback of our approach is that it ignores any portfolio effects in SARFT's decision-making. We do not allow the government to prefer a mix of movies, such as a certain number of comedies relative to action movies. These kinds of preferences would be difficult to estimate given our sample size, and more importantly, we are not aware of SARFT having such preferences.

8.2 Counterfactual Estimates

Our calculation of welfare is standard and follows papers such as Petrin (2002). For completeness, we present details in Appendix ???. Table 7 reports our results. We present our calculations on a per-year basis and use only the periods 2012-2014. While we also have data for 2015, we have data

²³Recall that movies with no U.S. production in the top 100 were always selected in the pre-2012 period. This approach mimics that outcome in our counterfactual calculations.

only for part of the year, which creates difficulty for our counterfactual calculations about which movies SARFT would keep or drop. The left panel shows the result in levels and the right panel shows percentage changes between the counterfactual and observed outcomes.²⁴ For instance, we observe an average of 638.13 million tickets sold per year over the three year period from 2012-2014. According to our main specification, Column 4 of Table 3, we calculate that if the 14 movies with the lowest admissions in China of the 34 were removed, there would be 625.96 million tickets sold. The right panel indicates that going from 625.96 to 638.13 million tickets sold is a 1.9% increase. In contrast, eliminating the top 14 movies by admissions from the set of 34 would reduce ticket sales to 549.04 million, and going from this number of ticket sales to the observed level would be an increase of 16.2%. Obviously, even among the 34 foreign movies with revenue-sharing contracts, there is a big difference between the top and bottom movies.

We are particularly interested in the results using our model of SARFT behavior. Whether we use parameters from the post-2012 or pre-2012 period (columns 1 and 3 in Table 6 respectively) makes little difference, so we focus on the post-2012 outcome. Although SARFT heavily weights revenue, SARFT's choices lead to substantially more impact than the *Bottom 14* column. The results from modeling SARFT's decision-making lead to annual sales of 605.43 million tickets, and going from 605.43 to 638.13 is a percentage increase of 5.4%.

The extra movies significantly impact the foreign and domestic share. In Table 7, the second row reports foreign movie admissions. The third row reports admissions to the non-excluded foreign movies when all movies are available. For instance, total admissions for foreign movies when all movies are available is 280.91 million. When the bottom 14 movies are eliminated, foreign admissions are 257 million. These remaining movies had 246.83 million admissions (the third row) when all movies were available. Thus, the percentage change for foreign movies that compete with the bottom 14 is $(246.83-257)/257=-3.96\%$, which appears in the third row of the right panel.

The SARFT model generates a percentage increase of foreign ticket sales of 24.3%, with foreign

²⁴We have also calculated standard errors for this table, but we do not report them to make the presentation more readable. They are small, particularly for the percentages, and similar across columns. The standard errors are below 11 for the admissions numbers and below 0.5 percentage points for the percentages. We use the delta method in order to calculate standard errors, which accounts for confidence intervals in the demand parameters. Currently, our standard errors for Table 7 do not account for estimation or simulation error in the SARFT probit calculations, which would affect the middle columns of each panel.

movies that compete with the newly introduced foreign movies experiencing a decline in ticket sales of 8.7%. That compares with domestic movies, which lose only 5.8% of tickets when the 14 foreign movies are added to the market. In this sense, foreign movies are a closer substitute amongst each other than with domestic movies.

This result is particularly striking when we consider the *Top 14* column. Going from the market without the top 14 movies to the *All Movies* column increases foreign ticket sales by 88.1%, with competing foreign movies experiencing a decline of 26%, as compared to a decline in domestic movies of only 10.6%. These large differences in the effect on foreign and domestic are driven by the standard deviation in the random coefficient on *Foreign* in Table 3. In this sense, it appears that foreign and domestic movies are not very close substitutes for each other. This result calls into question the value of import quotas as a way to protect the domestic movie industry, as it does appear the foreign movies greatly impact domestic movie-going.

We also compute the effect on consumer utils using the discrete choice model. Naturally, given that we effectively assume that price is constant, the change in utils closely tracks the change in quantities. However, Table 7 shows that the percentage change in utils is somewhat higher than the percentage change in quantities for the SARFT case, implying that it is particularly high-quality movies that are affected by liberalization. As discussed above, converting utils into dollar numbers is not straightforward as a result of a lack of price variation in this market. One result in the literature comes from de Roos & McKenzie (2014), which exploits the presence of discounted tickets on Tuesdays in Australia, and finds the own-price elasticity is about 2.5. Using a nested logit model in data from Hong Kong, Ho, Liang, Weinberg & Yan (2018) find an elasticity in the range of 5 to 6.5.²⁵ Neither of these papers account for consumption durability, so it is difficult to match them exactly. We calibrate our model to have a price coefficient of -0.126, which generates a product elasticity of 4. In that case, the movie market generates welfare of ¥20.8 billion. With such a large number, relatively small changes become important. Going from the restricted choice set implied by the post-2012 SARFT model to the observed choice set implies an increase in surplus of ¥1.172

²⁵Our understanding is that the 2.5 is a market elasticity whereas the 5-6.5 is a product elasticity, which potentially explains the discrepancy.

billion. While establishing the correct price coefficient for these calculations is not the focus of the paper, it seems clear that the results are economically meaningful.

We are also interested in comparing these results to what we find when using a static model that ignores consumption durability. i.e. the standard BLP results from Column 1 of Table 3. The results for this counterfactual are reported in the bottom panel of Table 7. We see that the static model generates larger results for welfare gains of liberalization than the dynamic model. Under the static model, counterfactual ticket sales are lower but ticket sales of foreign moves are slightly higher than under the dynamic model. But the biggest difference is in the domestic movies: the static model predicts that sales of domestic movies hardly changes at all, a 1.1% change when using the SARFT model.

The size of the deviation of the random coefficient on *Foreign* is similar in the static and dynamic results. Instead, the result appears to be due to dynamics: in the static model, consumers choose to see foreign movies every week, whereas the dynamic model predicts that consumers that prefer to see foreign movies still move on to domestic movies after seeing foreign movies. Keep in mind that the dynamic model could have matched the low substitutability in the static model by finding a much higher standard deviation in the random coefficient, whereas the static model cannot match the complex sequence of market shares predicted by the dynamic model. We find the low substitutability for the static model in Table 7 implausible, and we view this as evidence in favor of the dynamic model.

Table 7: Welfare and Market Share Effects of the Import Liberalization from 2012

		Exclusion based on				Percentage Change due to Exclusion			
		All movies	SARFT Model			Bottom 14	SARFT Model		
			Bottom 14	pre-2012	post-2012		Top 14	Bottom 14	pre-2012
Dynamic Model									
Annual Admissions	638.13	625.96	608.02	605.43	549.04	1.94%	4.96%	5.41%	16.23%
Annual Admissions of Foreign Movies	280.91	257.00	229.83	226.15	149.35	9.30%	22.28%	24.29%	88.09%
Annual Admissions of Competing Foreign Movies (when all movies are available)		246.83	211.41	206.42	110.51	-3.96%	-8.03%	-8.74%	-26.00%
Annual Admissions of Domestic Movies	357.22	368.96	378.19	379.28	399.68	-3.18%	-5.54%	-5.81%	-10.63%
Annual Consumer Welfare (Util)	2616.92	2549.37	2478.99	2469.3	2265.67	2.65%	5.57%	5.98%	15.50%
Static Model									
Annual Admissions	638.13	618.85	595.4	591.9	522.03	3.11%	7.18%	7.82%	22.24%
Annual Admissions of Foreign Movies	280.91	259.74	234.41	230.87	155.14	8.15%	19.89%	21.74%	81.07%
Annual Admissions of Competing Foreign Movies (when all movies are available)		246.83	211.41	206.42	110.51	-4.97%	-9.82%	-10.61%	-28.77%
Annual Admissions of Domestic Movies	357.22	359.11	360.99	361.03	366.89	-0.53%	-1.05%	-1.06%	-2.64%
Annual Consumer Welfare (Util)	933.55	894.38	844.26	836.42	685.89	4.38%	10.59%	11.63%	36.11%

The left panel presents the outcome in levels in millions, averaged across the three years from 2012-2014. The right side presents percentage changes of the left side. The column *All movies* is the observed outcome when 34 movie are selected for revenue-sharing. The rest present counterfactual restrictions to 20 movies. *Bottom 14* removes the lowest 14 movies by Chinese box office admissions, the *SARFT Model* columns make use of the Probit model estimated in Table 6 as described in Section 8.1. The *pre-2012* column uses Column 3 and the *post-2012* column uses Column 1. The *Top 14* column eliminates the top 14 movies by admissions. The percentage changes in the right panel are the percentage increase going from the restricted choice set to the *All Movies* choice set on the left panel. The top panel (Dynamic Model) uses the model with consumption durability based on Column 4 of Table 3. The bottom panel (Static Model) uses a standard static model and parameters from Column 1 of Table 3.

In these calculations, consumers have fewer choices in the counterfactual settings due to foreign movies being excluded. In a logit-based model such as ours, that implies that consumers draw less logit error terms, which can impact choices and the level of welfare for consumers.²⁶ In order to evaluate the importance of this issue, we recalculate the counterfactual calculations in Table 7 but this time, we replace lost movies in the consumer choice set so consumers always have six named movies available. In each period in which a movie is no longer available, we replace it with the next highest movie by admissions that week, typically the 7th highest level of admissions. In most cases, the movie has appeared in the top six at some time, so we have a movie fixed effect for that movie. If that is not available, we use the results of Table 4 to compute a movie fixed effect. We also use the results of the demand estimation with seven choices (Column 5 of Table 3) to compute the average decrease in utility in going from the 6th to 7th movie. More details about the replacement process appear in Appendix C.

Results appear in Table 8. While the changes in admissions and welfare move in the expected directions, the changes are quite small. For instance, going from having the top-14 movies unavailable with replacement to being available increases total admissions by 16.02%, whereas in the main specification, we saw that this number without replacement is 16.23%. This small difference between 16.23% to 16.02% is driven by the fact that the seventh movie in a market is much worse than a top 14 movie. We conclude that our results in Table 7 are not driven by the mechanics of counting logit error terms.

9 Conclusion

We study demand for movies in China. We propose a model that recognizes movies as *performance goods*: Choice sets rapidly evolve, consumers have limited time to devote to seeing movies in theaters, and consumers rarely want to see movies multiple times, which we term consumption durability. We propose a dynamic model of consumer demand that captures these features.

We apply the model to detailed administrative data on ticket sales drawn from a government

²⁶Ackerberg & Rysman (2005) address this issue in estimation by introducing a term that controls for the number of products in the choice set, motivated by a model of product crowding in unobserved utility space. However, we set the number of inside products to six throughout estimation, so this is not an issue for estimation, and we do not have the variation to estimate such a term.

Table 8: Welfare and Market Share Effects of the Import Liberalization from 2012 Allowing for Replacement of Excluded Movies

	Exclusion based on				
	All movies	SARFT Model			Top 14
		Bottom 14	pre-2012	post-2012	
Levels					
Annual Admissions	638.13	626.31	608.52	605.87	550.01
Annual Admissions of Foreign Movies	280.91	257.47	230.43	226.72	150.56
Annual Admissions of Competing Foreign Movies (when all movies are available)		246.83	211.41	206.42	110.51
Annual Admissions of Domestic Movies	357.22	368.85	378.09	379.15	399.46
Annual Consumer Welfare (Util)	2616.92	2551.03	2481.22	2471.09	2268.79
Percentage Change					
Annual Admissions		1.89%	4.87%	5.33%	16.02%
Annual Admissions of Foreign Movies		9.11%	21.97%	23.98%	86.58%
Annual Admissions of Competing Foreign Movies (when all movies are available)		-4.13%	-8.27%	-8.97%	-26.60%
Annual Admissions of Domestic Movies		-3.15%	-5.52%	-5.78%	-10.57%
Annual Consumer Welfare (Util)		2.58%	5.47%	5.91%	15.34%

This table is similar to Table 7 but the calculations replace movies that are not accepted into China in the consumer choice set, so consumers continue to have the same number of choices in the baseline and counterfactual settings.

agency. Like movie markets in other countries and other cultural goods products, ticket sales in China exhibit a stark decline in sales soon after their introduction. Whereas previous research used coefficients on age in static and reduced-form models to match this feature, we find that coefficients on age go essentially to zero when estimating with our model. Thus, it appears that consumption durability can well explain this feature of the data without relying on reduced-form age coefficients. We focus on a model with consumption durability and myopic consumers, which we show fits the data better than a forward-looking model in which consumers account for future movie releases.

We use the model to consider policy-relevant counterfactual scenarios. In particular, China effectively places a quota on the number of foreign movies that may be imported. This quota was increased from 20 to 34 in 2012. We evaluate the welfare increase from this change, and we find it to be significant, leading to a 6% increase in consumer welfare. Our results provide a measure of the consumer welfare cost of these types of quotas, hopefully to be accounted for in policies designed to protect domestic culture or local industries.

In addition, we find that there is relatively little substitution between foreign and domestic

movies. This result raises questions about the role of the quota as a tool for infant industry protection, as it appears that relaxing the quota would have a relatively low impact on domestic film production.

References

- Akerberg, D. A. & Rysman, M. (2005). Unobservable product differentiation in discrete choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics*, 36, 771–788.
- Aguiar, L. & Waldfoegel, J. (2018). Quality predictability and the welfare benefits from new products: Evidence from the digitization of recorded music. *Journal of Political Economy*, 126, 492–524.
- Barroso, A. & Llobet, G. (2012). Advertising and consumer awareness of new, differentiated products. *Journal of Marketing Research*, 49, 773–792.
- Berry, S. (1994). Estimating discrete choice models of product differentiation. *RAND Journal of Economics*, 25, 242–262.
- Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63, 841–890.
- Blonigen, B. A. & Soderbery, A. (2010). Measuring the benefits of product variety with an accurate variety set. *Journal of International Economics*, 82, 168–180.
- Broda, C. & Weinstein, D. E. (2006). Globalization and the gains from variety. *Quarterly Journal of Economics*, 121, 541–585.
- Bruno, H. A. & Vilcassim, N. J. (2008). Research note - Structural demand estimation with varying product availability. *Marketing Science*, 27, 1126–1131.
- Chu-Shore, J. (2010). Homogenization and specialization effects of international trade: Are cultural goods exceptional? *World Development*, 38, 37–47.
- Conlon, C. T. & Mortimer, J. H. (2013). Demand estimation under incomplete product availability. *American Economic Journal: Microeconomics*, 5, 1–30.
- Dalton, C. M., Gowrisankaran, G., & Town, R. (2019). Saliency, myopia and complex dynamic incentives: Evidence from medicare part d. *Review of Economic Studies*, In Press.
- Davis, P. (2006). Spatial competition in retail markets: Movie theaters. *RAND Journal of Economics*, 37, 964–982.
- De Groot, O. & Verboven, F. (2019). Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review*, 109, 2137–2172.
- de Roos, N. & McKenzie, J. (2014). Cheap Tuesdays and the demand for cinema. *International Journal of Industrial Organization*, 33, 93 – 109.
- Draganska, M. & Klapper, D. (2011). Choice set heterogeneity and the role of advertising: An analysis with micro and macro data. *Journal of Marketing Research*, 48, 653–669.
- Duan, N. (1983). Smearing estimate: A nonparametric retransformation method. *Journal of the American Statistical Association*, 78, :605–610.
- Einav, L. (2007). Seasonality in the U.S. motion picture industry. *The RAND Journal of Economics*, 38, 127–145.
- Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economics Review*, 84, 157–177.
- Ferreira, F., Petrin, A., & Waldfoegel, J. (2016). Preference externalities and the rise of China: Measuring their impact on consumers and producers in the global film market. Working paper, University of Minnesota.
- Ferson, W. E. & Constantinides, G. M. (1991). Habit persistence and durability in aggregate consumption: Empirical tests. *Journal of Financial Economics*, 29, 199–240.
- Francois, P. & van Ypersele, T. (2002). On the protection of cultural goods. *Journal of International Economics*, 56(2), 359–369.

- Gandhi, A. & Houde, J.-F. (2019). Measuring substitution patterns in differentiated-products industries. Unpublished manuscript, University of Wisconsin-Madison.
- Gentzkow, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, *97*, 713–744.
- Goeree, M. S. (2008). Limited information and advertising in the US personal computer industry. *Econometrica*, *76*, 1017–1074.
- Gowrisankaran, G. & Rysman, M. (2012). Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, *120*, 1173–1219.
- Greenwald, B. & Stiglitz, J. E. (2006). Helping infant economies grow: Foundations of trade policies for developing countries. *American Economic Review*, *96*, 141–146.
- Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, *50*, 1029–1054.
- Hanson, G. H. & Xiang, C. (2009). International trade in motion picture services. In *International Trade in Services and Intangibles in the Era of Globalization*, NBER Chapters (pp. 203–222). National Bureau of Economic Research, Inc.
- Hanson, G. H. & Xiang, C. (2011). Trade barriers and trade flows with product heterogeneity: An application to US motion picture exports. *Journal of International Economics*, *83*, 14–26.
- Hayashi, F. (1985). The permanent income hypothesis and consumption durability: Analysis based on Japanese panel data. *Quarterly Journal of Economics*, *100*, 1083–1113.
- Hendricks, K. & Sorensen, A. (2009). Information and the skewness of music sales. *Journal of Political Economy*, *117*, 324–369.
- Ho, J. Y., Liang, Y., Weinberg, C. B., & Yan, J. (2018). An empirical study of uniform and differential pricing in the movie theatrical market. *Journal of Marketing Research*, *55*, 414–431.
- Holloway, I. (2014). Foreign entry, quality, and cultural distance: Product-level evidence from US movie exports. *Review of World Economics*, *150*, 371–392.
- Honka, E., Hortaçsu, A., & Vitorino, M. A. (2017). Advertising, consumer awareness, and choice: evidence from the u.s. banking industry. *RAND Journal of Economics*, *48*, 611–646.
- Kokas, A. (2017). *Hollywood Made in China*. University of California Press.
- Krugman, P. R. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of International Economics*, *9*, 469–479.
- Kwak, J. & Zhang, L. (2011). Does China love Hollywood? An empirical study on the determinants of the box-office performance of the foreign films in China. *International Area Studies Review*, *14*, 115–140.
- Lee, F. (2006). Cultural discount and cross-culture predictability: Examining the box office performance of American movies in Hong Kong. *Journal of Media Economics*, *19*, 259–278.
- Lee, R. (2013). Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review*, *103*, 2960–3000.
- Leung, T. C. & Qi, S. (2020). Globalization and the rise of action movies in hollywood. Unpublished manuscript, Wake Forest University.
- Li, S., Mazur, J., Park, Y., Roberts, J., Sweeting, A., & Zhang, J. (2019). Repositioning and market power after airline mergers. Unpublished manuscript, University of Maryland.
- Magnac, T. & Thesmar, D. (2002). Identifying dynamic discrete decision processes. *Econometrica*, *70*, 801–816.
- Marvasti, A. & Canterbury, E. R. (2005). Cultural and other barriers to motion pictures trade. *Economic Inquiry*, *43*, 39–54.

- Maystre, N., Olivier, J., Thoenig, M., & Verdier, T. (2014). Product-based cultural change: Is the village global? *Journal of International Economics*, *92*, 212–230.
- McCalman, P. (2004). Foreign direct investment and intellectual property rights: Evidence from Hollywoods global distribution of movies and videos. *Journal of International Economics*, *62*, 107–123.
- Mehta, A., Rysman, M., & Simcoe, T. (2010). Identifying the age profile of patent citations. *Journal of Applied Econometrics*, *25*, 1179–1204.
- Moul, C. C. (2007). Measuring word of mouth’s impact on theatrical movie admissions. *Journal of Economics and Management Strategy*, *16*, 859–892.
- Moul, C. C. (2008). Retailer entry conditions and wholesaler conduct: The theatrical distribution of motion pictures. *International Journal of Industrial Organization*, *26*, 966–983.
- Musalem, A., Olivares, M., Bradlow, E. T., Terwiesch, C., & Corsten, D. (2010). Structural estimation of the effect of out-of-stocks. *Management Science*, *56*, 1180–1197.
- Orbach, B. Y. & Einav, L. (2007). Uniform prices for differentiated goods: The case of the movie-theater industry. *International Review of Law and Economics*, *27*(2), 129–153.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, *110*, 705–729.
- Rauch, J. E. & Trindade, V. (2009). Neckties in the tropics: A model of international trade and cultural diversity. *Canadian Journal of Economics*, *42*(3), 809–843.
- Sheu, G. (2014). Price, quality, and variety: Measuring the gains from trade in differentiated products. *American Economic Journal: Applied Economics*, *6*, 66–89.
- Sunada, M. (2012). Competition among movie theaters: An empirical investigation of the Toho-Subaru antitrust case. *Journal of Cultural Economics*, *36*, 179–206.
- Trajtenberg, M. (1989). The welfare analysis of product innovations, with an application to computed tomography scanners. *Journal of Political Economy*, *97*, 444–479.

Appendices

Appendix A Additional Tables

Table 9: Movie Characteristics for Movies That Ever Appear in the Top 6 for a Week

Variables	Unweighted			Admission-Weighted		
	(1) All	(2) Domestic	(3) Foreign	(4) All	(5) Domestic	(6) Foreign
Age (Week)	3.26	3.48	2.95	4.34	4.96	3.66
RunTime (Minute)	108.5	104.8	114.0	118.7	112.3	125.8
Indicator variables:						
IMAX	0.21	0.08	0.40	0.45	0.20	0.72
3D	0.34	0.21	0.54	0.51	0.32	0.72
Foreign	0.41	0	1	0.48	0	1
Action	0.38	0.25	0.58	0.51	0.33	0.71
Comedy	0.28	0.33	0.21	0.26	0.34	0.16
Drama	0.33	0.37	0.26	0.34	0.49	0.19

427 observations, 253 domestic and 174 foreign.

Table 10: Robustness Results for Table 3

Parameters		Differences	Perfect Fore- sight	Perfect Fore- sight	1-week fore- sight	Importance Sampling
Non-linear	Constant	9.189 (0.731)	10.526 (0.126)	10.425 (0.135)	9.645 (0.248)	14.634 (0.083)
	Enhanced(3D or IMAX)	1.437 (2.054)	1.531 (0.201)	1.524 (0.138)	1.509 (0.310)	1.203 (0.057)
	Foreign	2.821 (1.100)	3.343 (0.164)	3.316 (0.114)	3.038 (0.148)	5.476 (0.070)
	Discount Rate		0.000 (0.023)	0.500		
Linear	Age	-0.117 (0.632)	0.006 (0.081)	-0.190 (0.072)	-0.330 (0.072)	0.060 (0.089)
	Holiday	1.121 (0.616)	1.300 (0.413)	1.252 (0.395)	1.166 (0.405)	1.244 (0.464)
	Consumption Durability	Yes	Yes	Yes	Yes	Yes
	Micro-moments	Yes	Yes	Yes	Yes	Yes
	Forward Looking	No	Yes	Yes	Yes	No

1,464 observations. These models provide robustness checks for Table 3. The first column uses the first difference in ξ to form moments rather than the level. The second column allows for perfect foresight and estimates the discount rate to be zero. The third column imposes the discount rate $\lambda = 0.5$. The fourth column imposes $\lambda = 1$ and allows consumers to look only one period into the future. The fifth column is equivalent to Column 4 in Table 3 but uses importance sampling.

Appendix B Welfare Computation

For each year, we compare market shares and welfare from the observed choice set with 34 foreign movies to the welfare from a counterfactual choice set of 20 foreign movies, where the selection of the 20 movies is as described in the previous sub-section. In this subsection, we take the counterfactual choice set as given and define how to compute the resulting welfare change.

We denote the counterfactual choice sets as \tilde{C}_{gt} , $g = 1, \dots, G_t$. Some choice sets \tilde{C}_{gt} have the same set of movies as C_{gt} because they do not include any movies that have been excluded. But, for the choice sets C_{gt} that include an excluded movie, \tilde{C}_{gt} is a strict subset of C_{gt} . For the counterfactual set of choice sets, we employ the estimated mean utility and follow Equations 1-4 to compute the market share of each of the remaining movies week by week, solving for new choice probabilities and transitions.

To evaluate the welfare benefit of import liberalization on consumer welfare, we compute the welfare to consumers with and without the excluded movies, as follows:

$$\% \Delta CS = \int \frac{CS_i - \widetilde{CS}_i}{\widetilde{CS}_i} dF_i$$

where

$$CS_i = \sum_t \sum_{g \in G_t} s_{igt} \ln \left(1 + \sum_{j \in C_{gt}} e^{\delta_{jt} + \mu_{ijt}} \right)$$

$$\widetilde{CS}_i = \sum_t \sum_{g \in G_t} \widetilde{s}_{igt} \ln \left(1 + \sum_{j \in \widetilde{C}_{gt}} e^{\delta_{jt} + \mu_{ijt}} \right).$$
(7)

That is, we first compute the welfare for each consumer i facing a choice set \widetilde{C}_{gt} in each week t . Second, we sum up the welfare for each consumer i facing different choice sets according to her probability of facing each choice set, \widetilde{s}_{igt} . Third, we aggregate consumer welfare for each consumer i over all weeks to obtain \widetilde{CS}_i . Finally, we aggregate consumer welfare over all consumers. We compare the counterfactual consumer welfare to the consumer welfare from the observed data to compute the percentage change in consumer welfare. We compute total admissions and admissions to foreign and domestic movies in a similar way.

Appendix C Counterfactual calculations with replacement

In this section, we describe our method for calculating the case in which we assure consumers always have six movies in their potential choice set by replacing movies that are excluded by SARFT in our counterfactual analysis. In our main approach, as presented in Table 7, excluded movies are removed from the choice set, so consumers may have less than six named movies they can choose from. That is, C_t is reduced to five, four, or (for two weeks in our data) even three choices. In the computation presented in Table 8, we replace the excluded movies so that C_t has six elements in every period.

Our goal is to replace the excluded movie with the next best non-excluded movie available that period, which is typically the seventh most popular movie that week. In order to include the movie, we need the observable characteristics of the movie in order to calculate the interaction with random coefficients and we must take a position on what mean utility δ_{jt} the movie would have had if it had been in C_t that period. Mean utility consists of three elements, the movie fixed effect, the other explanatory variables (which are time-varying, such as the time trend, holiday weekend, and

month-of-year effects), and the unobserved quality ξ_{jt} .

For the movie fixed effect, if the movie ever appeared in \mathbf{C}_t in another period, we have estimated the movie fixed effect and we use that. If not, we use the prediction from the movie fixed effects regression (parameters are presented in Table 4), assuming the error term is zero. We apply the appropriate time effects for the period in question using estimates from the demand estimation (parameters are presented in Table 3), and we assume $\xi_{jt} = 0$. We denote this value $\delta_t^{(7)}$, the mean utility of the seventh most popular movie available that period. Similarly, we refer to the value of the sixth most popular movie as $\delta_t^{(6)}$. This value $\delta_t^{(6)}$ is computed as part of our estimation routine, so we treat it as observed in our counterfactual computations.

In addition, we want to be sure the newly added movie has lower mean utility than the existing movies, in particular that $\delta_t^{(7)} < \delta_t^{(6)}$. In order to do so, we utilize our estimation that allows for seven named choices (parameters presented in Column 5 of Table 3). We calculate the average difference between the mean utility of the sixth and seventh most popular movie. In particular, let $\tilde{\delta}_t^{(k)}$ be the mean utility of k th most popular movie in the specification with seven named choices per period. We calculate:

$$\Delta\delta = \frac{1}{T} \sum_{t=1}^T \left(\tilde{\delta}_t^{(7)} - \tilde{\delta}_t^{(6)} \right)$$

For the mean utility of the replacement movie, we use $\min\{\delta_t^{(7)}, \delta_t^{(6)} - \Delta\delta\}$. If we must replace two movies for that week, we use $\min\{\delta_t^{(8)}, \delta_t^{(6)} - 2\Delta\delta\}$ for the second replacement, and so on for the case of three replacements.

A final issue is that our model is not designed to handle movies that leave the set of six choices and then return. While one can imagine writing the model to address this, it came up for only one observation in our main specification and so we have not modeled this issue (see Footnote 15). This problem comes up in about 14% of weeks in this calculation however, as we are reaching into the 7th, 8th, or 9th most popular movies. When the replacement movie creates an adjacency issue, we simply keep going down the list to the 9th or 10th most popular movie to find one that does not create the adjacency issue.

Once we have filled in mean utilities for all of the replacement movies, the counterfactual calcu-

lations follow exactly as in the baseline case.