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

July 3, 2025

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 that is dynamic in the sense that long-lived consumers see movies only once and prefer not to go to the movies with very high frequency. 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 7.7% due to the import liberalization and that there is relatively little substitution between foreign and domestic movies. Revenue in China expands due to liberalization because of demand expansion, which provides insights into Chinese motivation for the quota.

Keywords: Demand Estimation, Choice Set, Trade Liberalization.

JEL classifications: L10, L82, F13

*We thanks many seminar audiences and especially our referees and editor for helpful comments.

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, how much this expansion led to substitution away from other movies, and the implications for revenue to the Chinese movie industry. We focus on distinguishing between the effects on foreign and domestic movies, which is relevant for evaluating industry protection policies.

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 music 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 than a reduced-form age profile.

In our model, consumers face an exogenously evolving choice set. Consumers have heterogeneous

[1] Not all cultural goods are performance goods, such as music on streaming services, and some performance goods may not be cultural goods, such as online courses in computer programming. Note that consumption durability has long been considered in the literatures on macroeconomic and finance to understand consumption dynamics (Hayashi, 1985; Ferson & Constantinides, 1991).

preferences over movie characteristics, and those preferences do not change over time. We assume consumers can see no more than one movie per week, and face reduced utility from going to the movies with high frequency, reflecting consumers' limited time for attending cinemas. 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 aggregate our model to the product level and find the level of unobserved quality for each movie-week that rationalizes the observed market share, and we 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 data well.

We apply our model to a dataset 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 movies experience an extreme drop-off in sales from week to week, our results show that the drop-off can be entirely explained by consumption durability. In particular, when we estimate a traditional static random coefficients logit model with a reduced-form age profile, we find that the age profile is strongly significant and negative, reflecting the steep drop-off 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.

We also find substantial heterogeneity in preferences for foreign movies, suggesting that foreign and domestic movies are not close substitutes. Substitution interacts with consumption durability in important ways. Consider a sequence of periods with the same movies available. Without consumption durability, consumers that prefer foreign movies will see them each period. But with consumption durability, consumers may exhaust the available foreign movies and then switch to domestic movies. That is, with consumption durability, consumers that see both foreign and domestic movies are not necessarily a sign of high substitution. We show how our model can generate complex substitution patterns that are not feasible under a static model.

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 are 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 7.71% due to the liberalization. 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 about the value of infant industry policies, as substitution between the foreign and domestic products is limited. Indeed, we find that losses to domestic film production from liberalization are more than offset by gains to domestic distributors, theaters, and government tax receipts. This result suggests that either the quota is motivated by cultural protection rather than economic protection or the Chinese government particularly values protection for domestic movie production as opposed to other parts of the supply chain. In addition, we find that if the consumption durability in preferences is ignored, the welfare benefit for consumers is underestimated, and the level of substitution between foreign and domestic movies is also underestimated.

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 affect foreign movie production in terms of genre and content (see for instance Leung & Qi, 2022). We do not address that issue here, although it may be important.

2 Literature

Our work contributes to literatures on trade in movies and cultural goods as well as demand estimation. We are related to the growing empirical literature on trade in motion pictures, such as Marvasti & Canterbery (2005), Hanson & Xiang (2011), Holloway (2014), and McCalman (2004), which study factors that lead US movies to be imported into a country, such as trade barriers and intellectual property policies. Relative to these papers, we focus on the import choices of a single country, China, and so focus on which movie characteristics determine entry. Ferreira, Petrin & Waldfogel (2016) estimate a structural model to evaluate the role of product quality in determining gains from trade in motion pictures. Similarly, we use a structural demand model to examine the welfare effects of import liberalization, although we highlight the importance of consumption durability.

Several papers study demand for movies, such as Einav (2007), which focuses on seasonality of demand, Moul (2007), which studies word-of-mouth in generating demand, de Roos & McKenzie (2014), which estimates a price elasticity by exploiting a discounting strategy. Particularly in China, Gil, Ho, Xu & Zhou (2024) and Chen, Yi & Yu (2024) estimate price elasticities to movies using more disaggregate data than ours. Accounting for consumption durability as we do is new in this literature, although we exploit the Chinese elasticity results to calibrate one part of our counterfactual. Hodgson & Sun (2025) studies the implications of vertical integration in the Chinese movie industry, although does so with a static demand system that does not address consumption durability.

Our paper adds to the lengthy literature on the benefits of new goods in industrial organization and trade. Whereas most industrial organization papers, such as Trajtenberg (1989) and Petrin (2002), study new goods in the context of innovation, we focus on the liberalization of a quota, which links to the literature on trade barriers. A leading trade 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; Sheu, 2014). Our paper particularly studies a cultural good, which are sometimes exempted from trade agreements. The protection of national culture 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).[2] In addition to protecting national culture, countries sometimes restrict the entry of foreign goods to encourage the development of local producers, so-called infant industry protection. See Greenwald & Stiglitz (2006) and the literature that follows. We show that benefits to liberalization were high in the Chinese movie context due to low substitution between foreign and domestic products, which limits the potential importance of infant industry protection.

Trade barriers are of 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.[3] Article IV of the GATT agreements in 1947 provides the conditions under which countries may impose quotas on foreign movies.[4]

Our paper is the first to estimate a dynamic model of movie demand or account for consumption durability in movies or other cultural goods. We build on the methodology developed by Berry, Levinsohn & Pakes (1995) to estimate a demand system for differentiated products with market-level data, and Gowrisankaran & Rysman (2012) to incorporate product durability and forward-looking consumers. Consumer dynamics arise in many contexts in addition to performance goods,

[2]Cultural goods and services "encompass values, identity and meanings that go beyond their strictly commercial value," according to a submission from the Coalition for Cultural Diversity to the UNESCO Convention on the Protection and Promotion of the Diversity of Cultural Expressions. Francois & 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.

[3]Marvasti & Canterbery (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.mpaa.org/wp-content/uploads/2015/03/MPAA-Theatrical-Market-Statistics-2014.pdf

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

such as durable goods, storable goods, and subscription goods (Gowrisankaran & Rysman, 2020). Our main specification that focuses on myopic consumers, but still with consumption durability, simplifies the problem of Gowrisankaran & Rysman (2012) substantially. 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 (2020), and De Groote & Verboven (2019).

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."(Rosen, 2002).

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 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, both of which are state-owned enterprises. There is no specific quota to import movies on a fixed-fee basis, and it is usually 20-30 per year. It does not appear that the prevalence of entering by fixed fee changed around the policy change in 2012, with the number of top movies entering by fixed fee remaining small.[5]

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, the same as domestic producers (about 30%), 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).

The way revenue from the box office is split between the various levels is determined by regulation and bargaining. The government gets 8.3% in tax revenues. Of the remaining, 25% go to the foreign producer of a revenue-sharing movie, whereas for domestic and co-produced movies, that number is 30%. The rest goes to the distributors and theaters. The share going to theaters varies very little, if at all, across foreign and domestic movies.[6] Thus, from the perspective of receipts in China, the difference between a foreign revenue-sharing movie and a domestic movie is that China retains about

[5]In 2010-2012, the number of movies entering China via flat fee among the top 100 movies by North American box office that year were 4, 8, and 7. In 2014 and 2015, the numbers were 6 and 7. The outlier is 2013, when 11 of the top 100 movies entered by fixed fee. That year had some unusual circumstances. The movie *Ironman 3* was aiming for coproduction status but was denied at a late date, and *Hunger Games: Catching Fire* also faced a distributional issue. It is difficult to say why the fixed-fee market is not more active given the quota on revenue sharing, although we note that we are not aware of fixed fees between distributors and producers being used anywhere in the world. Perhaps it is difficult to contract on fixed fees when products have highly uncertain outcomes.

[6]See *China-International Film Co-Production Handbook* produced by the Motion Picture Association and the China Film Co-Production Corporation.

77% $(1-(1-0.083)\times 0.25)$) of box office revenue from foreign movies but 100% of box office revenue for domestic movies. Thus, if liberalization decreases domestic admissions by less than 77% of the increase in foreign admissions, China gains overall from liberalization even ignoring any benefits to consumers. We return to this calculation in our counterfactual results in Section 9.

While we can calculate how much revenue accrues to the levels of taxes, production, distribution, and theatrical showings, determining the exact gains and losses from liberalization (and the political economy of determining the quota) is complex because producers own some theaters and furthermore, the government owns the main distributors and may have ownership stakes at other levels. We do not account for these ownership patterns, but from the perspective of infant industry protection, even if direct losses to domestic producers from competition through liberalization are offset by gains from ownership stakes at other levels of the supply chain (such as cinemas), those gains will not incentivize investment in domestic production. To the extent that industry protection is aimed at domestic production as opposed to domestic producers in all of their lines of business, our approach is sufficient.

All films, whether they are foreign, domestic, or coproduced, face censorship by SARFT. 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 are not allowed in any 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."

Similar to other markets, price variation in the Chinese movie market is limited. While prices vary by time of day, day of the 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.[7] Conditional on theater and time, prices do not vary by whether a movie is foreign or domestic. Similar to Einav (2007) and others in this literature, we do not attempt to estimate a price coefficient. Rather, we capture the mean level of movie utility with a movie fixed effect and present counterfactual results

[7] Orbach & 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.

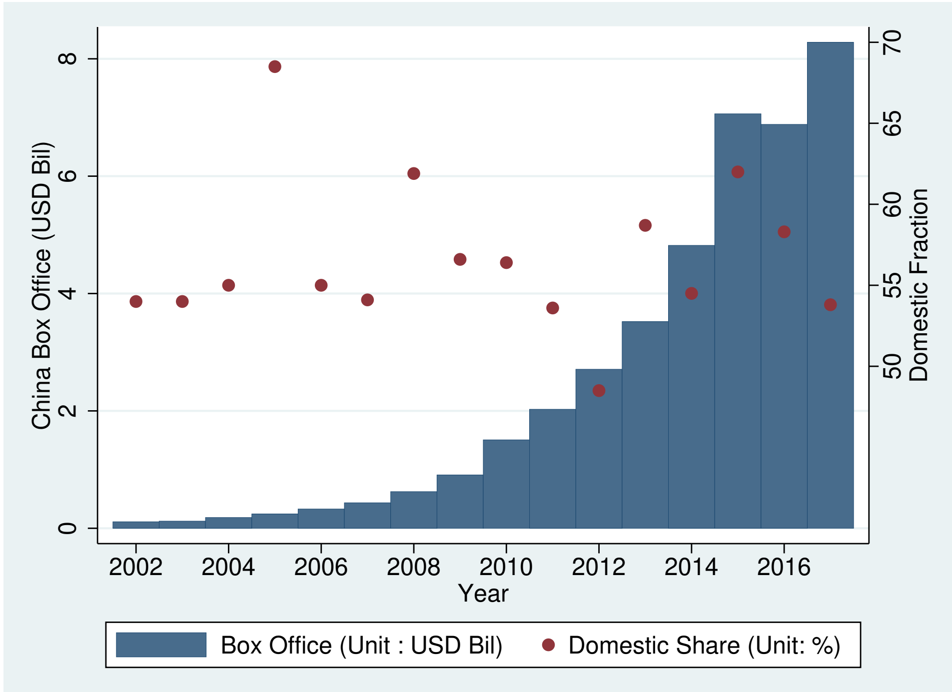

Figure 1: Chinese box office revenue and domestic Share.

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.

Figure 1 depicts that the share of domestic movies at the box office remained at about 55% over a long period, which may relate to the import restrictions by China on foreign movies.[8] The 55% share is higher than those in European countries documented in Hanson & Xiang (2009). 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.

4 Data

The empirical analysis is based on a novel dataset from a media consulting firm that works with 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 each week. Beginning in January 2012,

[8]The data for this figure were collected by the authors from several online sources, particularly reports by Entgroup.

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. The number of tickets sold is also referred to as *admissions*. 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.

4.1 Descriptive statistics

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% 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 of admissions going to each rank of movie, i.e. the top-ranked movie, the second-ranked movie, and so on. We average this over the 183 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.

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

[9] Lee (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

	Unweighted			Admission-Weighted		
	(1)	(2)	(3)	(4)	(5)	(6)
Variables	All	Domestic	Foreign	All	Domestic	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.

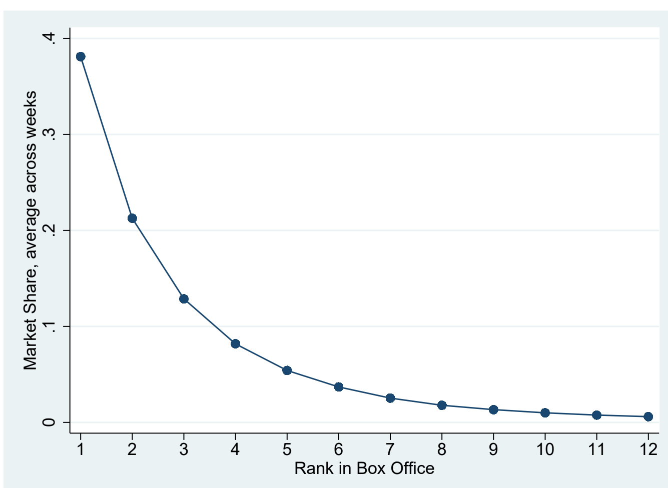

Figure 2: Average share of ticket sales by weekly sales rank

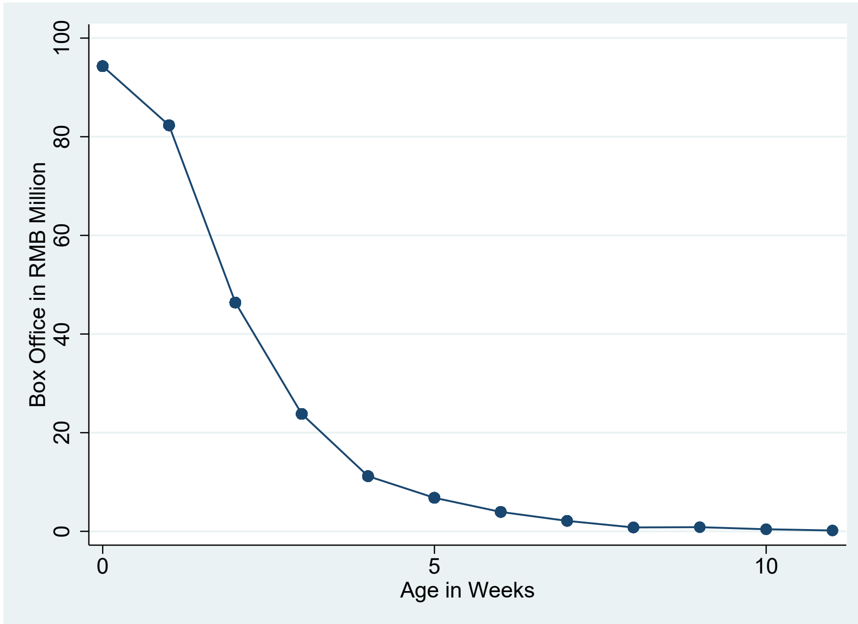

Figure 3: Age profile of movie revenue

of the log of admissions 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]

4.2 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

[10] In 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).

[11] Note 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 a holiday.

We further include a linear time trend in the month as well as year fixed effects to capture the dramatic increase in the Chinese movie market documented in Figure 1. 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-of-the-year dummies is meant to capture anecdotal evidence that SARFT's treatment of foreign movies varies by season.

4.3 Market size and choices

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 a 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 consist 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 increase 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.

There are two or three foreign movies in the top 6 in more than half the weeks, and no foreign movies available in only 7 of the 183 weeks in our data. Note that throughout the paper, we treat

[12] For 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

Note: 1,464 observations, 8 movies in each week, 183 weeks.

co-produced movies as domestic and we combine fixed fee and revenue sharing movies as foreign movies. Among movies that make it into the top six, fixed fee movies account for about 5.4% of total box office revenue relative to 41.6% for foreign movies overall so breaking them out leads to very small sample sizes. 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 A1 in Appendix A.

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 the 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. In the mathematical model below, this size is denoted M. The population data is obtained from the China Statistical Yearbook. To compute market shares, we divide admissions (recall that admissions is quantity, not revenue) of movie j in week t by the market size. Let q_{jt} be the admissions of movie j in week t. Then, $s_{jt}^{\texttt{data}} = q_{jt}/M$ 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%.

4.4 Data for micro-moments

We collect the annual frequency of movie watching from *China Moviegoer Survey Report 2012-2013*, a survey conducted by a Chinese consulting firm called Entgroup. The survey was conducted online in February and March of 2013. The 6,027 respondents are consumers who watched at least one movie in the theater in the previous year. Entgroup's sampling method ensures the geographical distribution of survey respondents is consistent with the empirical geographic distribution of movie-goers in 2012. 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 6 movies in the previous year.[13]

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: rapid 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

[13]We can use the survey data to get perspective on the total movie-going. In our admissions data, admissions from February 2012 to January 2013 is 463.82 million. As an approximation, suppose 23.2% see two movies, 19.2% see five movies, and 57.5% see 9 movies. That would imply 70.2 million households see a movie in that year, about 19.2% of our potential market size of 354 million. The annual growth rate in admissions is 34.4%, and naturally, we want to allow for consumers to see a movie in one year and not another. For both reasons, we need a potential market size substantially higher than 70.2 million. Thus, our choice of the potential market size is not particularly large by the standards of the discrete choice demand estimation literature.

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 allow consumers to experience disutility from seeing a movie soon (within a few weeks) after seeing another movie, which we interpret as another facet of limited time, but could also be seen as a form of consumption durability.

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 relax this assumption 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 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 of 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,

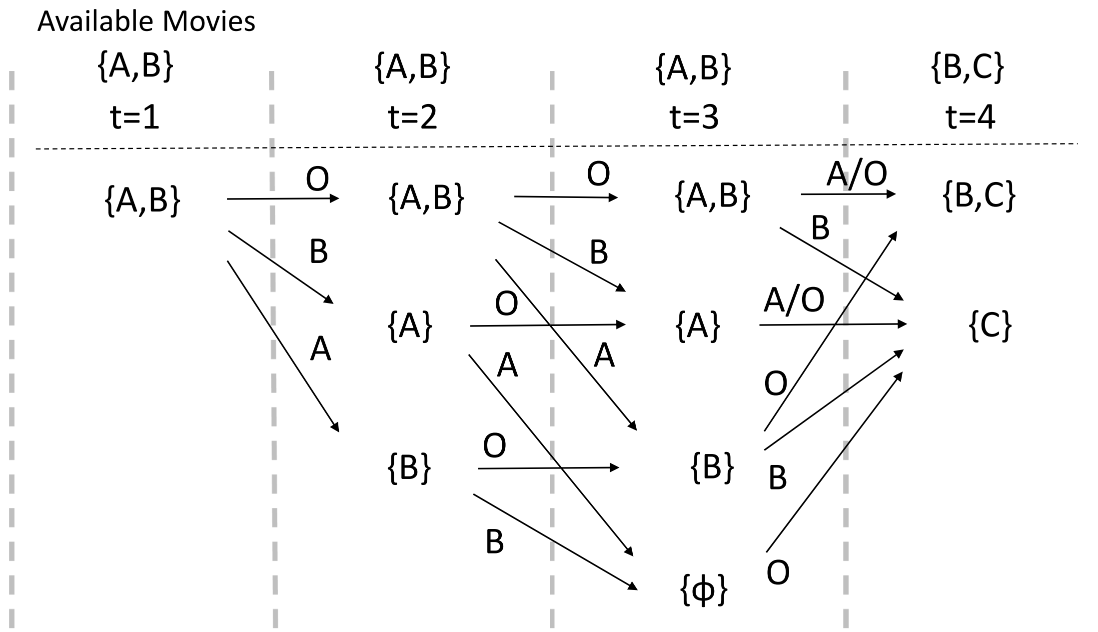

Figure 4: Simplified representation of the demand model

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 M indexed by i face discrete, finite time indexed by $t = 1, \ldots, T$, where one period represents one week in the data. The set of all movies ever available can be indexed by j from 1 to J. These are the named movies that a consumer can see only once. 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}, $g = 1, \ldots, G_t$. 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.

We augment each choice set C_{gt} with three additional choices: the option not to purchase ($j = 0$), a generic foreign movie ($j = J + 1$), and 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.

Denote the choice that consumer i makes in period t as $d_{it} \in \{0, \ldots, J + 2\}$. Denote the history of choices by i up to period t as H_{it}, so $H_{it} = \{d_{it-1}, d_{it-2}, \ldots, d_{i1}\}$. Let the function $C(H_{it}, \mathbf{C_t})$ return consumer i's choice set in t:

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

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 adds the outside options to the choice set. It must be that $C\left(H_{it}, \mathbf{C_t}\right)$ equals an element C_{gt} of $\mathcal{C}_t$.

Because consumers may not want to go to the movies week after week, we allow consumer utility to depend on how recently the consumer has seen any movie. Let $w(H_{it})$ be the number of weeks since consumer i has seen a movie: $w(H_{it}) = \min\{\tau : \tau \in \mathbb{Z}_+, d_{i,t-\tau} > 0\}$, with $w(H_{it}) = \infty$ if i

has never seen a movie (i.e. $d_{i\tau} = 0 \; \forall \tau = 1, \ldots, t-1$). Here, $\mathbb{Z}_+$ are the strictly positive integers. Note that $w(\cdot)$ does not distinguish whether the last movie seen was named or generic, and so treats them symmetrically.

Let the utility to consumer i from choosing movie j in period t with history H_{it} be denoted by $u_{ijt}(H_{it})$. Consumer i solves:

$$\max_{j \in C(H_{it}, \boldsymbol{C}_t)} u_{ijt}(H_{it}).$$

We assume that utility takes on the functional form:

$$u_{ijt}(H_{it}) = x_{jt}\beta + \xi_{jt} + \mu_{ijt} + \rho \mathbb{1}\{w(H_{it}) < 5\} + \varepsilon_{ijt}.$$

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 plays the role of the econometric error term. The parameter ρ measures the *recency effect* of having seen a movie, so we expect to find $\rho < 0$. This allows consumers who have seen a movie in the last four weeks to be less likely to see a movie in the current week. We also experimented with different numbers of weeks, and allowing the recency effect to decline over the four weeks. 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's match to the product based on observable characteristics. Following Berry (1994) and Berry et al. (1995), we specify it as:

$$\mu_{ijt} = \sum_{l=1}^{L} x_{jlt}\sigma_l \nu_{il}$$

where $\nu_{il} \sim \mathcal{N}(0,1)$. Thus, ν_{il} 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. These preferences are assumed to be constant over time for each consumer. The parameters β, ρ, and σ_l, $l = 1, \ldots, L$ are to be estimated. We refer to estimation parameters together as $\theta = \{\beta, \rho, \{\sigma_l\}_{l=1,\ldots,L}\}$. Furthermore, for convenience, we denote the mean utility of product j in period t as $\delta_{jt} = x_{jt}\beta + \xi_{jt}$. We assume $\delta_{0t} = 0$ and $x_{0t} = 0$ for all t.

5.3 Market shares

Given these assumptions, history affects consumer choices through two channels, the recency effect and in determining the choice set. Conditional on knowing the choice set and the last period a consumer purchased, there is no further value to knowing history, so we define choice probabilities on these variables. The conditional probability $P_{ijt}(C(H_{it}, \boldsymbol{C}_t), w(H_{it}))$, the probability of i choosing j in t conditional on having choice set $C(H_{it}, \boldsymbol{C}_t)$ and have last purchased $w(H_{it})$ periods ago, is:

$$P_{ijt}(C(H_{it}, \boldsymbol{C}_t), w(H_{it})) = \begin{matrix} \dfrac{\exp\left(\delta_{jt} + \mu_{ijt} + \rho \mathbb{1}\{w(H_{it}) < 5\}\right)}{\sum_{k \in C(H_{it}, \boldsymbol{C}_t)} \exp\left(\delta_{kt} + \mu_{ikt} + \rho \mathbb{1}\{w(H_{it}) < 5\}\right)} & j \in C(H_{it}, \boldsymbol{C}_t) \\ 0 & \texttt{otherwise} \end{matrix} . \tag{1}$$

In Figure 4, $P_{ijt}(C(H_{it}, \boldsymbol{C}_t), w(H_{it}))$ 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 $\mathbf{C_t}$:

$$B_{gg't} = \left\{ j : C_{g't+1} = \underbrace{(C_{gt} \backslash \{j\})}_{\texttt{current}} \cup \underbrace{(\mathbf{C_{t+1}} \backslash \mathbf{C_t})}_{\texttt{entering}} \backslash \underbrace{(\mathbf{C_t} \backslash \mathbf{C_{t+1}})}_{\texttt{exiting}} \right\}. \tag{2}$$

To track consumers over time, we must track the recency of their consumption in addition to their choice sets. Let the variable $w \in \{1, 2, 3, 4, 5\}$ represent the number of periods since last seeing a movie, where $w = 5$ indicates not having seen a movie in the last four periods. Let s_{igwt} be the share of consumers of type i (that is, with draw ν_i) with choice set g in period t who saw a movie w periods ago. Thus, $\sum_{g=1}^{G_t} \sum_{w=1}^{5} s_{igwt} = 1$. We refer to s_{igwt} as the unconditional probability or unconditional share. To compute s_{igwt}, we assume that there is only one possible choice set in the first period: $\mathcal{C}_1 = \{\mathbf{C_1}\}$ and that no consumer has seen a movie in four periods. Thus, $s_{i151} = 1$ for all i and $s_{igw1} = 0$ for all g and $w \neq 5$. Unconditional shares evolve as follows:

$$
\begin{aligned}
s_{ig'1t+1} &= \sum_{g=1}^{G_t} \sum_{w=1}^{5} \sum_{j \in B_{gg't}, j \neq 0} P_{ijt}(C_{gt}, w) s_{igwt} \\
s_{ig'wt+1} &= \sum_{g=1}^{G_t} \mathbb{1}\{0 \in B_{gg't}\} P_{i0t}(C_{gt}, w-1) s_{igw-1t} \quad w \in \{2,3,4\} \\
s_{ig'5t+1} &= \sum_{g=1}^{G_t} \mathbb{1}\{0 \in B_{gg't}\} \left(P_{i0t}(C_{gt}, 4) s_{ig4t} + P_{i0t}(C_{gt}, 5) s_{ig5t} \right) \\
&\quad \forall g' = 1, \ldots, G_{t+1}, t = 2, \ldots, T-1.
\end{aligned} \tag{3}
$$

The first equation in (3) defines the share of consumers transitioning to choice set g' and $w = 1$ in $t+1$, so these are consumers that purchase in t. The second equation defines the share going to g' and w from 2 to 4. Naturally, this requires that the consumer does not purchase in choice t and so chooses $j = 0$. The last equation defines the share for g' and $w = 5$. This is similar to the second equation but includes consumers with two values of w in t: $w = 4$ and $w = 5$. That is, both consumers that purchased four weeks ago and purchased more than four weeks ago end up with no recency effect ($w = 5$) in period $t+1$ if they do not purchase in t. The "∀" line at the bottom of (3) applies to all three lines.

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

$$
\widehat{s}_{jt} = \int \sum_{g=1}^{G_t} \sum_{w=1}^{5} P_{ijt}(C_{gt}, w) s_{igwt} f(\boldsymbol{\nu}_i) d\boldsymbol{\nu}_i \tag{4}
$$

where $f(\boldsymbol{\nu}_i)$ is the distribution of consumer types $\boldsymbol{\nu}_i$, and $\boldsymbol{\nu}_i$ is the $L \times 1$ vector of elements ν_{il}.

5.4 Forward-looking behavior

In some performance goods settings, our assumption of myopic behavior may not be reasonable.[14] In this subsection, 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.

[14] For 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.

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 in period t with history H_{it} to be:

$$V_{it}(H_{it}) = \ln\left(\sum_{j\in C(H_{it},\boldsymbol{C}_t)} \exp\left(\delta_{jt} + \mu_{ijt} + \rho\mathbb{1}\{w(H_{it}) < 5\} + \lambda V_{it+1}(H_{it+1})\right)\right)$$

where the consumer's history in $t+1$ is the choice j in t added to H_{it}, that is, $H_{it+1} = \{j, H_{it}\}$. The variable λ is the discount rate. For this calculation, we assume that $V_{iT+1}(H_{iT+1}) = 0$ for all i and H_{iT+1}. Thus, for a consumer in the final period T, the choice problem is the same whether we use myopic or forward-looking behavior.

We can define the utility to i from movie j as:

$$u_{ijt}(H_{it}) = \delta_{jt} + \mu_{ijt} + \rho\mathbb{1}\{w(H_{it}) < 5\} + \lambda V_{it+1}(H_{it+1}) + \varepsilon_{ijt}$$

We use this utility specification to adjust Equation 1 accordingly. The rest of the model, such as the determination of s_{igt}, remains the same. We estimate this model by backward induction. For a given guess of the parameters, we calculate the utility and probability of each choice in the last period. We 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. To identify the discount rate, we rely on variation in the continuation value that is not reflected in the current values, as in Magnac & Thesmar (2002). This kind of variation is natural in perfect foresight models. 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 in future weeks. Our prior belief is that consumers heavily discount this continuation value, and indeed, that is our finding 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. An alternative approach based on limited information assumptions would be 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 be no more difficult to estimate but for the extra parameter. 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.

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 Reduced-form results

There are several implications of our model of consumption durability relative to a static model of demand. One is that a movie introduced at the same time as a strong rival is likely to have relatively higher sales later in its life cycle. In the opening week, consumers choose between the two movies and many will choose the strong rival. But in later weeks, the rival is no longer part of those consumers' choice sets, and those consumers are more likely to choose other movies. In contrast, a static model with an age profile predicts that a movie introduced at the same time as a strong rival will have lower sales throughout its lifetime.

Consumption durability also has complex implications for substitution patterns. Consumers that strongly prefer foreign movies may still go to domestic movies if they have exhausted all of the foreign movies available. whereas a static model allows those consumers to go to foreign movies every period. Another implication is that large expansions of the choice set may be more beneficial under consumption durability because consumers cannot repeatedly consume the top products. However, recognizing that consumers have limited time to consume performance goods may moderate the benefits of choice expansion. There are more implications one may consider, but they are outside the focus of our study, such as the implications for release date strategies.

In this section, we provide reduced-form analysis to show that our data exhibit the effects that our model highlights. We focus on the first hypothesis, that consumption durability implies relatively high sales later in the life cycle for movies released alongside popular movies. We also provide a reduced-form estimate of the extent of substitution between foreign and domestic movies, but delay analysis of the more complex timing of substitution, as well as welfare effects, for the structural estimation.

In order to explore the first hypothesis, we define a set of top movies, for instance, the top twenty movies over our sample by Chinese box office admissions. We define a dummy variable that equals one for movie j if movie j's first week coincides with week one or two of a top movie. We then interact the dummy variable with indicators for movie j being in week two or later. We interpret this interaction as an indicator for consumption durability affecting movie j. Formally:

$$\texttt{ConsumptionDurability}_{jt} = \mathbb{1}\{j \texttt{ is introduced alongside a top movie}\}\mathbb{1}\{\texttt{age}_{jt} > 1\}$$

We estimate linear models of the log of box office sales as a function of our indicator for consumption durability along with movie fixed effects, age fixed effects, and calendar control variables. We use the same calendar controls we use in our structural estimation: year fixed effects, a month time trend, month-of-the-year fixed effects, and a dummy for being a holiday weekend. We estimate via OLS. To best mimic our structural estimation, our sample consists of the top six movies each week by Chinese box office. We consider several definitions of top movies, such as the top 20 and top 30 over the sample, or the top six or ten each year.

Results appear in Table 3. The coefficient ranges from 0.235 to 0.351. In three of the four cases, the parameter is significant at a 95% level of significance, with the t-statistic for the fourth case equal to 1.958. The coefficients imply a percentage change increase in sales ranging from 27% to 42% across specifications, although admittedly on a relatively small base. Interestingly, the coefficient increases as we more tightly restrict the set of top movies, suggesting that the effect grows as the popularity of the rival movie grows. However, the standard error on consumption durability also grows, and differences across specifications are not statistically significant. Overall however, we take these results as evidence in favor of consumption durability.

A potential concern for this regression is that we take the set of rival movies as uncorrelated with the age profile so, for instance, distributors do not introduce movies that distributors expect to have flatter age profiles at the same time as popular movies. We test for this possibility directly in our structural estimation below. We do not find evidence for heterogeneous age profiles.

A second finding in the results of our structural model below is the relatively low substitution between foreign and domestic movies. In order to explore this feature in reduced form, we construct

Table 3: Consumption durability

Number of top movies across all years	20	30
Consumption durability indicator	0.351	0.295
	(0.134)	(0.125)
R-squared	0.868	0.867

Number of top movies across each year	6	10
Total Number of Top movies	21	35
Consumption durability indicator	0.301	0.235
	(0.142)	(0.120)
R-squared	0.867	0.867

Notes: 1,071 observations in each regression. Standard errors in parenthesis. Dependent variable is log sales in a week for a movie. Consumption durability is a dummy for being in week 2 or later and being introduced at the same time as a top movie. All specifications include movie fixed effects, age fixed effects by week, holiday weekend fixed effects, month-of-the-year fixed effects, year fixed effects, and a month time trend.

Table 4: Substitution between foreign and domestic movies

Ln(sales of foreign movies)	Monthly		Weekly	
Domestic	-0.290	-0.420	-0.528	-0.637
	(0.105)	(0.174)	(0.105)	(0.070)
Holiday	No	Yes	No	Yes
Monthly Trend	No	Yes	No	Yes
Month-of the-Year FE	No	Yes	No	Yes
Year FE	No	Yes	No	Yes
Observation	42	42	182	182

total sales of all foreign and all domestic movies by week and regress the log of foreign movie sales on the log of domestic sales. A coefficient higher than -1 suggests imperfect substitution. In our most complete specification, we include a holiday week dummy, a month time trend, month-of-the-year fixed effects, and year fixed effects. Results are in Table 4. While the coefficient is always negative and significantly different from zero, it is always statistically significantly higher than -1 as well. We find a coefficient of -0.528 when we do not include controls and -0.637 when we include all of our controls. We also estimate at the monthly level rather than the weekly level. Standard errors are larger, but the coefficients are still significantly above -1.[15]

A natural next question is how consumption durability and foreign-domestic substitution interact.

[15] A potential concern might be weather shocks that affect overall moviegoing and create correlation between foreign and domestic admissions, driving the coefficient away from -1. However, weather patterns are varied across China, and most likely "average out." For example, Shanghai and Beijing are 1000km apart, and experience both their highest average rainfall and highest average temperatures in different months.

For this more complicated question, we turn to the results of the structural model.

7 Structural estimation

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

7.1 Aggregate data moments

We cannot compute Equation 4 analytically. For this step, we use simulation. We draw S values of ν_{il}^s, $l = 1, \ldots, L$ and $s = 1, \ldots, S$.[16] For a given set of parameters θ and a guess of mean utilities δ_{jt}, we compute $P_{ijt}(C_{gt}, w)$ and then s_{igwt} for each C_{gt}, w, 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 $\widehat{s}_{jt}(\theta, \boldsymbol{\delta})$, where $\boldsymbol{\delta}$ is the vector of elements δ_{jt}.

As in Berry et al. (1995), we recover $\boldsymbol{\delta}$ for any set of parameters θ via the fixed point equation:

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

As above, $s_{jt}^{\texttt{data}}$ 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.[17]

[16]Denoting 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.

[17]The Berry (1994) proof of the invertibility of utility and market shares relies on the gross substitutability of products. See also Berry, Gandhi & Haile (2013). Gross substitutability does not necessarily hold in dynamic models. For instance, raising utility for one product in one period in our model raises sales of that product, which implies those consumers do not have the product in their choice set in the next period and are thus *more* likely to purchase another product, which violates substitutability.

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. As is typical in dynamic models, there is an initial conditions problem of determining the distribution of consumer choice sets at the start of the data. We assume that in the first period of the data, no consumers have seen a movie. In order to minimize the impact of this assumption on our results, we drop the first four weeks of data in forming our moments. Four weeks should be sufficient because there is frequent turnover in which movies are available, so 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 8 list 1,432 observations rather than 1,464.[18]

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. Our base specification places random coefficients on three variables: the constant term, an indicator for whether a movie is foreign, and an indicator for whether a movie was 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, particularly adding a random coefficient on *Action*. The generic domestic and foreign outside options each have their own fixed effect for quality and are subject to the time-varying explanatory variables (time trend, holiday, and month-of-year effects). The foreign generic movie (choice $J+1$) is also affected by the random coefficient on *foreign*.

We assume that all explanatory variables are exogenous. Recall that price is not an explanatory variable. However, the presence of consumer heterogeneity terms (σ_l) means we still need additional moments to achieve identification. Our first set of additional moments follows Berry et al. (1995). Because we take product introductions as exogenous, we use sums over the characteristics of other

[18] A 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 dataset. One movie leaves the top six in week 75 and returns in week 76. We keep this movie in $\boldsymbol{C}_t$ for $t = 75$, and drop the sixth most popular movie from $\boldsymbol{C}_{75}$.

movies in the top six in the same week as exogenous variables to construct additional moments. 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 (2020) recommend instruments that emphasize how differentiated a product is from others on the market. We construct these for the characteristics 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 is interacted with whether the movie in question is enhanced, and so this exogenous variable is high only for enhanced movies. We multiply the interaction variable by the econometric error term to construct additional moments.

7.2 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 23.2%, the probability of watching 4-6 movies is 19.2%, with the rest watching 7 or more.

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 and how recently they have seen a movie but also how many times they have been to the movies. That is, we denote the state of a consumer as $\{C_{gt}, w(H_{it}), n(H_{it})\}$ where $n(H_{it})$ is the number of movies that i has seen in the previous year, up to seven. We define $n(H_{it}) = \min\{\sum_{\tau=0}^{51} \mathbb{1}\{d_{it-\tau} > 0\}, 7\}$. We track the population that has seen different numbers of movies similar to how we track the weeks since the last movie in Equation 3. 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.

To be clear, this new state variable does not affect consumer decision-making. The consumer still cares only about the choice set and the time since seeing a movie. Rather, tracking the number of movies seen 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 $n = 0, 1, .., 6, 7$. 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 $P_{i4-6} = \sum_{n=4}^{6} P_{in} / \sum_{n=1}^{7} P_{in}$. Recall that P_{i7} refers to the probability i saw 7 or more movies. 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:

$$E[m_2(\theta)] \quad = \quad E \begin{bmatrix} P_{1-3}^{\mathtt{data}} - P_{1-3}(\theta) \\ P_{4-6}^{\mathtt{data}} - P_{4-6}(\theta) \end{bmatrix} \quad = \quad 0 \ , \tag{5}$$

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

$$E[m(\theta)] \quad = \quad E \begin{bmatrix} m_1(\theta) \\ m_2(\theta) \end{bmatrix} \quad = \quad 0 \ . \tag{6}$$

Here, $m_1(\theta)$ are the aggregate-data moments as discussed in Section 7.1. 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, we compute the weighting matrix following Conlon & Gortmaker (2025), particularly Case C of their Appendix E.

7.3 Identification

We finish this section with a brief heuristic discussion of how variation in data guides the results we find. Our model includes the same parameters as Berry et al. (1995) and thus, we can think of parameter results in a similar way. Products with different characteristics attract different levels of market share depending on the presence of alternative products. To the extent that 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 the entrance of a popular foreign movie, observing high market shares for 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. 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). While we attempt to provide reduced-form analysis of these effects in Section 6, an advantage of a structural model in a setting like this is that it resolves these complex interactions in a coherent and parsimonious way.

Focusing specifically on the random coefficient on the constant term, this parameter is identified without micromoments by the extent of substitution between the outside good and the inside goods, such as because of calendar effects. In addition, the parameter is heavily affected by the micromoments. They show extensive repeat moviegoing, which the model matches by setting a high standard deviation on the random coefficient on the constant term. Repeat purchase is moderated by a highly negative recency coefficient, which also causes reductions in the overall movie-going share in the weeks following a popular release.

8 Empirical results

This section discusses the empirical results obtained from the demand model described in the previous section.

8.1 Parameter results and consumption durability

We present parameter results from estimating our model in Table 5. Column 1 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 require a strong reduced-form age profile to match the kind of declines in market share that we see in the data (as evidenced in Figure 3). Column 2 adds the micro-moment. The micromoment contains more repeat purchase behavior than is generated by the static model. Under this specification, we see a much higher standard deviation of the random coefficient on the constant term in order to generate consumers that go to the movies many times. The age trend has an even larger negative coefficient. Because high heterogeneity means that movie-lovers would see a popular movie every period, a more negative age trend is necessary to match the age profile observed in the data.

Table 5: Demand estimates

		1	2	3	4	5	6
Heterogeneity	Constant	1.404	14.875	1.394	10.38	7.204	19.314
		(1.147)	(1.920)	(1.369)	(1.064)	(2.719)	(2.097)
	Enhanced(3D or IMAX)	0.186	0.076	0.965	1.547	0.072	4.884
		(3.075)	(0.419)	(2.528)	(1.405)	(3.074)	(0.858)
	Foreign	3.356	3.877	4.069	3.242	1.116	3.926
		(0.649)	(0.386)	(1.059)	(0.982)	(0.419)	(1.040)
	Recency (ρ)					-9.503	-7.671
						(4.677)	(1.514)
Linear	Age	-0.587	-0.85	-0.341	-0.016	-0.509	-0.019
		(0.039)	(0.063)	(0.094)	(0.089)	(0.094)	(0.128)
	Holiday	0.333	1.747	0.377	1.149	0.533	0.939
		(0.117)	(0.276)	(0.118)	(0.183)	(0.143)	(0.203)
	Consumption Durability	No	No	Yes	Yes	Yes	Yes
	Micro-moments	No	Yes	No	Yes	No	Yes

Specifications include movie fixed effects, month-of-year fixed effects separately for foreign and domestic movies, year fixed effects, and a month time trend. *Consumption Durability* refers to whether consumers can see movies only once (Yes) or the model is static (No). 1,432 observations. 427 movies. 179 weeks.

Column 3 estimates our dynamic model without the *recency* term. This is the model with myopic consumers who experience consumption durability but with $\rho = 0$. Adding consumption durability causes the coefficient on *age* to drop almost in half. Column 4 adds the micro-moment to Column 3. This specification leads to dramatic increases in the standard deviation of the random coefficient on *enhanced* and even more so on the constant term. That is, the way to match the repeat viewing in the survey data is to increase consumer heterogeneity so some consumers greatly 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.[19]

[19]The age trend is not separately identified from movie fixed effects and a week time trend, which is one reason we

Columns 5 and 6 add the *recency* term, with and without micro-moments. We find that ρ, the parameter on *recency*, is negative and statistically significant. The pattern is the same as in Columns 3 and 4: Adding micro-moments increases heterogeneity and causes the age profile to be close to zero and statistically insignificant. We take Column 6 to be our main results going forward. A feature of Column 6 is that the standard deviation parameters on the variables *enhanced* and *foreign* are statistically significant and economically large. This will drive our result in the next subsection that there is relatively muted substitution between foreign and domestic movies.

In thinking about identification, note that consumption durability does not necessarily imply that the *age* coefficient would be zero. Consumption durability implies that demand falls with age, 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 the coefficient on *age* driven to zero.

One way to evaluate the importance of consumption durability is to ask how often consumers face a restricted choice set. If we consider consumer types (that is, draws ν_i) that expect to see at least one movie in a year, we find that on average, they face a restricted choice set in 29.65% of weeks. That is, in about three out of ten weeks, movie-going consumers have already seen one of the movies available. Naturally, consumption durability interacts with consumer heterogeneity, so durability is even more important for consumers that particularly value movies. Conditioning on types that see on average two or more movies in a year, consumers face a restricted choice set 41.77% of weeks. We do this calculation for the year February 2012 to January 2013, which matches our micro-moment sample frame. Overall, we find that consumption durability plays a role in a significant percentage of consumer decisions.[20]

8.2 Substitution patterns

We are interested in evaluating substitution patterns, both intertemporally and between foreign and domestic movies. To do so, we consider a decrease in utils of 0.738 for a single movie, which

use a month trend. The results are robust to alternative treatments of the calendar time effects, such as using only year dummies.

[20]We also regress movie fixed effects on movie characteristics. Discussion and results appears in Appendix B and Table B1.

corresponds to a 10 RMB price increase if the price coefficient is -0.0738. We motivate this coefficient choice below, but it calibrates our model to reasonable price elasticities as measured by other papers. We change the utility of the top domestic movie at the midpoint of our dataset, which we number as movie 2207. This movie ranked second in weeks 94 and 95, fourth in week 96, and then dropped out of the top 6 thereafter. We change the utility of movie 2207 in week 94 only. We calculate the percentage change in ticket sales for movie 2207, for all other domestic movies (including the generic domestic option), and for all foreign movies (including the generic foreign option) for the week of the price change and the four weeks following. Results appear in Table 6.

Table 6: Substitution pattern for domestic movie

Week	Movie of Price Change	Other Domestic Movies	Foreign Movies
94	-38.98%	14.38	5.33
95	13.48	-5.35	-2.98
96	12.64	-1.5	-0.73
97		0.29	0.11
98		0.11	0.07
Long-run	-22.68	11.63	2.1

This table reports percentage quantity changes from a decrease in utils of movie 2207, a domestic movie, in week 94 (its first week on the market) to approximate a 10RMB price increase. The price increase lasts one period. The sales rank of the movie is 2 in weeks 94 and 95, 4 in week 96, and is out of the top 6 thereafter. The *Other domestic movies* and *Foreign movies* columns include the generic domestic and foreign options, respectively. The row *Long-run* sums the change in sales over the weeks 94-98 and divides by the quantity from week 94.

Focusing on the first column, the price change leads to a 39% decline in quantity in the period of the price change. However, that leads to a 13.5% increase in the next period, and a similarly large increase (over a relatively small base) in the following period. We are also interested in comparing the short and long-run elasticity. To calculate the long-run change, we sum the total change in quantity over five weeks and divide by the quantity from week 94. In this way, the denominators from the week 94 change and the long-run change are the same. The last row of the table reports the long-run change, and it is only -22.7%, less than 60% of the within-period change. Thus, as expected, the long-run elasticity is less than the short-run elasticity as consumers adjust by forgoing the movie and seeing it in later periods. Keep in mind that a model with only a reduced-form age profile and no consumption durability would predict no change in quantity in any periods following the price change.

We are also interested in substitution patterns to other movies. Table 6 shows that the price change leads to substantially more substitution to other domestic movies than foreign movies, with domestic movies increasing in the concurrent period by 14.4% and foreign movies increasing by only 5.3%. This difference is striking because the baseline sales of other domestic movies in this period is 2.78 million and foreign movies is 4.85 million, about 75% higher. That is, substitution is higher to domestic movies even though they have a substantially lower market share, which highlights the importance of the random coefficient approach to estimation.

Another way to evaluate substitution is with the diversion ratio (see Conlon & Mortimer, 2021, for a discussion of diversion ratios). Under the price increase, the number of tickets sold that period for movie 2207 goes from 2.18 million to 1.33 million, a decrease of 0.85 million. Other domestic movies increase from 2.78 to 3.18 million, an increase of 0.4 million. Thus, 47% (i.e., 0.4/0.85) of the loss to movie 2207 goes to other domestic movies. In contrast, foreign movie admissions increase from 4.85 to 5.11 million, leading to a diversion ratio of 30.6%, and 22.4% going to no purchase.

Strikingly, market shares for rival movies decline in week 95. In this sense, movies are intertemporal complements rather than substitutes. This is in part because consumers that switch to other movies in week 94 can no longer see those movies in week 95. Quantity changes switch sign over the five weeks, although at very low magnitudes, as the implications of the price change and durability work their way through the different options. As with the own-price effect, the long-run cross-price effects are lower than the short-run effects.

The effect is similar if we consider a price change to a foreign movie. We consider a change in price of the same level to movie 1323, an enhanced foreign movie, which appears in week 89 at number 1, and is then 4th, then 6th, and then drops out of the top 6. Results appear in Table 7. Results are similar, with a long-run own-price change less than the short-run, more substitution to foreign movies than domestic movies, and shifting signs over time.

8.3 Robustness

To establish robustness, we consider several alternative models. Robustness results appear in Table A2 in Appendix A. First, foreign action movies are particularly popular, so we should control

Table 7: Substitution pattern for foreign movie

Week	Movie of Price Change	Domestic Movies	Other Foreign Movies
89	-40.71%	4.77	11.58
90	11.62	-0.71	-2.6
91	4.21	-0.04	0.4
92		0.01	0.24
93		0	0.12
Long-run	-30.68	3.26	9.73

This table reports percentage quantity changes from a decrease in utils of movie 1323, a foreign movie, in week 89 (its first week on the market) to approximate a 10RMB price increase. The price increase lasts one period. The sales rank of the movie is 1 in week 89, 4 in week 90, 6 in week 91, and is out of the top 6 thereafter. The *Domestic movies* and *Other foreign movies* columns include the generic domestic and foreign options, respectively. The row *Long-run* sums the change in sales over the weeks 89-93 and divides by the quantity from week 89.

for action movies before we reach conclusions about the importance of foreign versus domestic movies. We re-estimate our main demand model (Column 6 of Table 5) allowing for a random coefficient on *Action*, the dummy for action movies. Results appear in Column 1 of Table A2. The standard deviation of the random coefficient on *Action* is statistically insignificant, and the coefficient on foreign hardly changes. Other parameters are similar and the age profile coefficient is still very close to zero and statistically insignificant.

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 2 of Table A2. These results appear very similar to our main specification. That is not surprising given the low market shares associated with low-ranked movies.

Another possible concern is that we construct our moments based on the assumption that $E[\xi_{jt}|Z_{jt}] = 0$. 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 Column 3 of Table A2.

We also consider a model in which consumers have perfect foresight as to what movies will be available, as described in Section 5.4. The specification is otherwise identical to that in Column 6 of Table 5. We perform a grid search over values of the discount rate λ from 0 to 1 and estimate the remaining parameters by non-linear optimization. We find that the objective function is minimized at a discount rate of $\lambda = 0$.[21] 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 the main specification, but standard errors change because of the extra parameter.[22] This result appears in Column 4 of Table A2.

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 0.5, we find the age profile coefficient increases in magnitude to -0.182, although is still statistically insignificant. 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 with λ fixed at 0.5 appears in Column 5 of Table A2.

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 Column 6 of Table A2 (see also Footnote 16).

Another possibility is that our assumption that all movies have the same age profile is restrictive. We might imagine that different types of movies have different age profiles. We focus on more popular movies versus less popular movies. We rerun the main specification (as in Column 6, Table 5) but with *age* interacted with whether a movie was ever ranked number 1 for a week (which is 108 of our 427 movies). In unreported results, we find that the coefficients on both *age* and *age* interacted with the top movie indicator are insignificantly different from zero.

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

[22] We find a relatively large standard error for λ of 0.431. 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. Interestingly, eliminating the recency term (setting $\rho = 0$) also leads to estimating $\lambda = 0$ but with a much lower standard deviation.

9 Counterfactual experiments

Since 2012, China has agreed to increase the import quota for foreign movies from 20 to 34 per year. 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.

9.1 Selecting movies under counterfactual quotas

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

9.1.1 A model of movie selection by SARFT

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, treating a movie as an observation, on which movies 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 Boxofficemojo.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. For robustness, we estimate two separate regressions, one for before and one for 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:

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

where $\eta \sim \mathcal{N}(0,1)$ and we observe $Y = 1$ if the movie is selected for revenue sharing by SARFT. In constructing the top 100, we do not consider movies that enter China 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 8, which reports the share of top movies (ordered by North American box office) selected for revenue-sharing by SARFT.[23] 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.[24]

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.[25] We also 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-U.S., 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 dataset in category 2 and all of them receive revenue-sharing. We drop these movies from our probit regression, so the

[23] Note 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.

[24] For movies with Chinese release dates close to their U.S. release dates, North American revenue would be unknown to 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.

[25] We do not attempt to measure other factors that SARFT appears to account for, such as whether the movie glorifies a foreign military, or is about religion or the occult. These are difficult to quantify. We briefly explored a machine-learning approach based on movie reviews, but with only a limited number of observations in each regression, and with much of the outcome explained by box office revenue, it did not appear that we had a dataset suitable for such an approach.

Table 8: 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.

earlier period has 395 rather than 400 observations. As we discuss below, we account for this feature in how we compute counterfactual outcomes.

9.1.2 Results of the SARFT model

The marginal effects from estimating the probit model appear in Table 9. 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 the box office by 1% is to increase the probability of selection by 0.2 percentage points.[26]

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

[26] We 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 $Y = 0$.

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 the 10th percentile of box office revenue. We provide two measures of fit. The first is the correlation coefficient between the dependent variable (the indicator for SARFT selection for revenue sharing) and the model's predicted probability for revenue sharing if $\eta = 0$. This statistic has a lower bound at zero, which would occur if we estimated the probit with only a constant term. The statistic has a maximum value of 1. We find values above 0.73 for all specifications. We also report the Pseudo R^2 of McFadden (1974). It also ranges from 0 to 1, and we find it is always above 0.5.

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.

9.1.3 Constructing SARFT's selection under counterfactual quotas

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 at least six non-U.S. movies. In fact, we observe fewer than six movies without the involvement of

Table 9: Probit model of SARFT's selection of movies for revenue-sharing

	Period 2012-15		Period 2008-2011	
ln(Box Office)	0.212	3.701	0.075	2.167
	(0.039)	(1.347)	(0.028)	(0.897)
ln(Box Office)2		-0.095		-0.057
		(0.04)		(0.024)
US and non-US producer	0.047	0.048	0.078	0.068
	(0.052)	(0.048)	(0.038)	(0.034)
Chinese producer involved	0.307	0.333	0.138	0.093
	(0.210)	(0.216)	(0.270)	(0.218)
No US producer	0.082	0.066		
	(0.282)	(0.276)		
Rated R	-0.229	-0.215	-0.156	-0.137
	(0.042)	(0.040)	(0.030)	(0.031)
IMAX	0.145	0.146	0.097	0.098
	(0.076)	(0.073)	(0.062)	(0.059)
3D	0.256	0.259	0.330	0.306
	(0.071)	(0.070)	(0.101)	(0.100)
In(RunTime)	0.308	0.284	0.406	0.356
	(0.178)	(0.163)	(0.137)	(0.127)
Action	0.162	0.153	0.180	0.157
	(0.07)	(0.06)	(0.05)	(0.051)
Comedy	-0.029	-0.028	-0.056	-0.051
	(0.06)	(0.06)	(0.03)	(0.029)
Drama	-0.058	-0.045	-0.091	-0.077
	(0.06)	(0.05)	(0.04)	(0.032)
2009			-0.001	0.002
			(0.06)	(0.051)
2010			-0.066	-0.059
			(0.05)	(0.041)
2011			-0.077	-0.072
			(0.04)	(0.039)
2013	-0.042	-0.034		
	(0.07)	(0.06)		
2014	-0.031	-0.037		
	(0.07)	(0.06)		
2015	-0.043	-0.033		
	(0.07)	(0.06)		
$\rho(Y_i, \hat{P}_i)$	0.732	0.745	0.745	0.746
Pseudo R^2	0.498	0.510	0.536	0.549
Observations	398	398	395	395

The sample is 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. Parameters are marginal effects and standard errors are in parentheses. 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. The line $\rho(Y_i, \hat{P}_i)$ reports the correlation coefficient between the SARFT's observed decision (the dependent variable) and the predicted probability of the decision from the model. The next line reports *Pseudo* R^2 based on McFadden (1974).

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.

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 9. 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 our counterfactuals, we calculate outcomes for each of the 100 simulations and report averages. 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 (that is, the quantity of tickets sold) in North America. 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.

9.2 Counterfactual estimates

Table 10 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 for only 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.[27] 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 6 of Table 5, we calculate that if the 14 movies with the lowest admissions of the 34 were removed, there would be 627.95 million tickets sold. The right panel indicates that going from 627.95 to 638.13 million tickets sold is a 1.62% increase. In contrast, eliminating the top 14 movies by admissions from the set of 34 would reduce ticket sales to 573.15 million, and going from this number of ticket sales to the observed level would be an increase of 11.34%. Obviously, even among the 34 foreign movies with revenue-sharing contracts, there is a large 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 9 respectively) makes little difference, so we focus on the post-2012 outcome. SARFT heavily weights revenue so the impact of SARFT's choices is closer to the *Bottom 14* column than the *Top 14*. The results from modeling SARFT's decision-making lead to annual sales of 612.37 million tickets, so liberalization represents a percentage increase of 4.21%.

The extra movies significantly impact the foreign and domestic share. In Table 10, 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 256.36 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 (246.83-256.36)/256.36=-3.72%, which appears in the third row of the right panel.

The SARFT model generates a percentage increase of foreign ticket sales of 22.84%, with foreign movies that compete with the newly introduced foreign movies experiencing a decline in ticket sales of 9.8%. That compares with domestic movies, which lose only 6.87% of tickets when the 14 foreign

[27] We have also calculated standard errors for this table, but we do not report them to make the presentation more readable. Standard errors 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 10 do not account for estimation or simulation error in the SARFT probit calculations, which would affect the middle columns of each panel.

movies are added to the market. In this sense, foreign movies are closer substitutes to each other than to 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 73.15%, with competing foreign movies experiencing a decline of 31.88%, as compared to a decline in domestic movies of only 13.07%. These large differences echo the substitution patterns we discussed in Section 8.2. 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 not appear that the foreign movies greatly impact domestic movie-going.

We are also interested in computing the dollar or yuan value of the change in consumer welfare. As discussed above, converting utility into a currency-valued number is not straightforward as a result of a lack of price variation in this market. We calibrate our price coefficient to match results in the literature. We focus on Gil et al. (2024), which studies moviegoing in China in a dataset related to ours during the same time period, and finds a product elasticity of 1.74. It is similar to Chen et al. (2024), which also studies moviegoing in China and finds an elasticity of 2.0. Calibrating our model to have a product elasticity of 1.74 leads to a price coefficient of -0.0738.[28] Using the post-2012 SARFT model of movie availability, we find that liberalization increases surplus by ¥1.185 billion (about US$300 million in 2015). Details on how we calculate the price coefficient and the yuan-valued welfare change are in Appendix C, but we generally follow papers such as Petrin (2002). Note that none of the existing papers with price elasticities address consumption durability, so they are not perfect for our needs. But while establishing the correct price coefficient for these calculations is not the focus of our paper, it seems clear that our results are economically meaningful.[29]

We display the percentage change in consumer welfare in Table 10. We find that liberalization

[28]There are other papers that calculate price elasticity for movies. de Roos & McKenzie (2014) exploit 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. We focus on the results in China with related data.

[29]Mechanically, the price coefficient linearly scales the welfare increase we find. For instance, if we had found a price coefficient of -0.1467 instead of -0.0738 (i.e., double what we found), we would find a welfare increase of 3.86% instead of 7.71% (i.e., half of what we found) The relationship between the price coefficient we find and the elasticity we calibrate to is more complex, but still monotonic. That is, higher elasticities lead to higher (in absolute value) price coefficients.

under the post-2012 SARFT model leads to an increase in utility of 7.71%.[30] Table 10 shows that the percentage change in utility is somewhat higher than the percentage change in quantities for the SARFT case, 7.71% relative to 4.21% in the post-2012 specification, implying that it is particularly high-quality movies that are affected by liberalization.

In addition to the increase in consumer welfare, we can also use our results to evaluate the contribution to revenue for the Chinese government and Chinese firms. Recall from Section 3 that the difference on this issue between a domestic and a revenue-sharing foreign movie is that whereas 100% of the revenue of a domestic movie stays in China, only 77% of revenue from a foreign movie stays in China (with the revenue going to tax receipts, distributors, and theaters). Thus, the Chinese movie industry as a whole gains from liberalization if domestic admissions decrease by less than 77% of the increase in foreign movies. Using the post-2012 SARF policy column, we see that liberalization leads domestic admissions to fall from 383.58 to 357.22 million, a decrease of 26.36, whereas foreign movies admissions increase from 228.79 to 280.91, an increase of 52.12. The increase in foreign admissions is 97% more than the decrease in domestic sales, another manifestation of the low substitution between foreign and domestic movies. Assuming that movie prices are similar across foreign and domestic films, our result suggests that Chinese revenue benefited from liberalization. This is consistent with the quota being driven either by cultural protection rather than just economic protection or because the Chinese government puts a high weight on protecting domestic movie production relative to other parts of the industry.[31]

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 5. The results for this counterfactual are reported in the bottom panel of Table 10. We see that the static model generates a lower welfare gain from liberalization than the dynamic model, 5.73% relative to 7.71%. Under the static model, counterfactual ticket sales are lower, but ticket sales of foreign movies are similar to the dynamic model. The biggest difference is in the domestic movies: the

[30] In this calculation, we display the ratio of the dollar-valued change in surplus to total industry revenue rather than starting or ending surplus. If we used surplus in the denominator, we would have to address the integrating constant in calculating consumer surplus. See Appendix C. Using observed industry revenue in the denominator addresses this issue. We do not display the levels (the left panel) because they are subject to the integrating constant.

[31] A fuller accounting of the benefits to China from liberalization would include Chinese investment in foreign movie production, which we do not address.

static model predicts that sales of domestic movies hardly change at all, less than a 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 10 implausible, and we view this as evidence in favor of the dynamic model.

Table 10: Welfare and market share effects of the import liberalization from 2012

	Quantities					Percentage Increase from Liberalization			
			SARFT Model				SARFT Model		
Dynamic Model with Recency	All movies	Bottom 14	Pre-2012	Post-2012	Top 14	Bottom 14	Pre-2012	Post-2012	Top 14
Annual Admissions	638.13	627.95	614.34	612.37	573.15	1.62	3.87	4.21	11.34
Annual Admissions of Foreign Movies	280.91	256.36	231.92	228.79	162.24	9.58	21.17	22.84	73.15
Annual Admissions of Competing Foreign Movies		246.83	211.41	206.42	110.51	-3.72	-8.86	-9.8	-31.88
Annual Admissions of Domestic Movies	357.22	371.59	382.43	383.58	410.91	-3.87	-6.59	-6.87	-13.07
Annual Consumer Welfare (Util)						3.66	7.17	7.71	19.19
Static Model									
Annual Admissions	638.13	618.72	595.15	591.65	520.99	3.14	7.23	7.86	22.48
Annual Admissions of Foreign Movies	280.91	260.03	235.02	231.51	156.41	8.03	19.58	21.41	79.6
Annual Admissions of Competing Foreign Movies		246.83	211.41	206.42	110.51	-5.08	-10.06	-10.85	-29.34
Annual Admissions of Domestic Movies	357.22	358.7	360.13	360.15	364.57	-0.41	-0.81	-0.81	-2.02
Annual Consumer Welfare (Util)						2.33	5.28	5.73	14.56

The left panel presents the outcome in levels in millions, averaged across the three years from 2012-2014. The column *All movies* is the observed outcome when 34 movies 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 9 as described in Section 9.1. 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. *Annual Admissions of Competing Foreign Movies* is the quantity sold for the remaining movies when all movies are available (i.e, under liberalization). It is from raw data and the same across panels. The top panel (Dynamic Model) uses the model with consumption durability based on Column 6 of Table 5. The bottom panel (Static Model) uses a standard static model and parameters from Column 1 of Table 5. We do not report standard errors for readability sake but they are below 11 for all items in the left panel and below 0.05 for all items in the right panel.

As a robustness check, we also consider these counterfactuals under the model without recency (based on Column 4 of Table 5). We present this result in the Appendix in Table A3. In this case, we find consumer welfare increases somewhat more, 8.88% relative to 7.71%, with other results scaled up similarly. As expected, the recency term moderates demand.

An issue with our calculations is that they assume consumers have fewer choices in the counterfactual settings due to foreign movies being excluded. In a logit-based model such as ours, this implies that consumers draw fewer logit error terms, which can impact choices and the level of welfare for consumers.[32] In order to evaluate the importance of this issue, we recalculate the counterfactual calculations, 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 seventh 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 B1 to compute a movie fixed effect. We also use the results of the demand estimation with seven choices (Column 2 of Table A2) to compute the average decrease in utility in going from the sixth to the seventh movie. More details about the replacement process appear in Appendix D.

Table A3 presents results. The results are in the bottom panel, and we do the calculations with the model without recency, so it is comparable to the top panel of the table. 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 15.72%, whereas without replacement, the number is 16.32%. The change is small because the seventh movie in a market is much worse than a top 14 movie and including a low-value movie does not impact results. We conclude that our results are not driven by the mechanics of counting logit error terms.

[32] Ackerberg & 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.

10 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 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 consumer welfare increase from this change, and we find it to be significant, leading to a 7.71% increase in consumer welfare. Further, we find that the Chinese government and firms benefit from liberalization because losses to Chinese movie production are more than offset by gains to tax receipts, distribution, and cinemas. This result is driven by our finding low substitution between foreign and domestic movies. Our paper provides a measure of the economic cost of these types of quotas, hopefully to be accounted for in policies designed to protect domestic culture or local industries.

Our results are explained in part because we find that there is relatively little substitution between foreign and domestic movies. Low substitution raises questions about the role of the quota as a tool for infant industry protection, as we find that relaxing the quota would have a relatively low impact on domestic film production. Consumption durability appears to be an important factor in this market, and likely related markets such as books and theater performances. Future work exploring the implications of consumption durability in other contexts, such as for release strategy, appears valuable.

References

Ackerberg, 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.

Berry, S. (1994). Estimating discrete choice models of product differentiation. *RAND Journal of Economics*, *25*, 242–262.

Berry, S., Gandhi, A., & Haile, P. A. (2013). Connected substitutes and invertibility of demand. *Econometrica*, *5*, 2087–2111.

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.

Chen, L., Yi, L. X., & Yu, C. (2024). The welfare effects of vertial integration in China's movie industry. *American Economic Journal: Microeconomics*, *16*, 204–235.

Chu-Shore, J. (2010). Homogenization and specialization effects of international trade: Are cultural goods exceptional? *World Development*, *38*, 37–47.

Conlon, C. & Gortmaker, J. (2025). Incorporating micro data into differentiated products demand estimation with PyBLP. *Journal of Econometrics*, In Press.

Conlon, C. & Mortimer, J. H. (2021). Empirical properties of diversion ratios. *RAND Journal of Economics*, *52*, 693–726.

Dalton, C. M., Gowrisankaran, G., & Town, R. J. (2020). Salience, myopia, and complex dynamic incentives: Evidence from Medicare Part D. *The Review of Economic Studies*, *87*, 822–869.

De Groote, 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.

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., & Waldfogel, 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. (2020). Measuring substitution patterns in differentiated-products industries. Unpublished manuscript, University of Wisconsin-Madison.

Gil, R., Ho, C.-Y., Xu, L., & Zhou, Y. (2024). Vertical integration and market foreclosure in media markets: Evidence from the Chinese motion picture industry. *Journal of Law and Economics*, *67*, 143–193.

Gowrisankaran, G. & Rysman, M. (2012). Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, *120*, 1173–1219.

Gowrisankaran, G. & Rysman, M. (2020). A framework for empirical models of dynamic demand. Unpublished manuscript, Boston University.

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.

Hodgson, C. & Sun, S. (2025). Heterogeneity in vertical foreclosure: Evidence from the Chinese film industry. Unpublished manuscript, Yale University.

Holloway, I. (2014). Foreign entry, quality, and cultural distance: Product-level evidence from US movie exports. *Review of World Economics*, *150*, 371–392.

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. (2022). Globalization and the rise of action movies in Hollywood. *Journal of Cultural Economics*, In Press.

Magnac, T. & Thesmar, D. (2002). Identifying dynamic discrete decision processes. *Econometrica*, *70*, 801–816.

Marvasti, A. & Canterbery, 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 Hollywood's global distribution of movies and videos. *Journal of International Economics*, *62*, 107–123.

McFadden, D. L. (1974). Conditional logit analysis of qualititative choice behavior. In P. Zarembka (Ed.), *Frontiers in Econometrics* (pp. 105–142). New York: Academic Press.

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.

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.

Rosen, S. (2002). The wolf at the door: Hollywood and thefilm market in china. In E. J. Heikkila & R. Pizarro (Eds.), *Southern California and theWorld* (pp. 49–77). Westport, CT: Praeger.

Sheu, G. (2014). Price, quality, and variety: Measuring the gains from trade in differentiated products. *American Economic Journal: Applied Economics*, *6*, 66–89.

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 A1: Movie characteristics for movies that ever appear in the Top 6 for a week

	Unweighted			Admission-Weighted		
	(1)	(2)	(3)	(4)	(5)	(6)
Variables	All	Domestic	Foreign	All	Domestic	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 A2: Robustness results for demand estimation (Table 5)

		1	2	3	4	5	6
Parameters		Action	7 named movies	First Differences	Perfect Foresight	Perfect Foresight	Importance Sampling
Heterogeneity	Constant	19.48	17.914	17.5	19.314	19.686	16.159
		(2.115)	(1.743)	(1.874)	(2.724)	(2.135)	(2.628)
	Enhanced(3D or IMAX)	4.993	4.565	4.482	4.884	4.984	1.932
		(0.995)	(0.861)	(0.902)	(0.939)	(1.136)	(0.983)
	Foreign	4.081	3.624	3.453	3.926	3.904	5.253
		(1.023)	(1.040)	(1.034)	(1.061)	(1.211)	(0.650)
	Action	0.422					
		(1.684)					
	Recent	-7.66	-6.974	-6.871	-7.671	-7.607	
		(1.653)	(1.300)	(1.338)	(1.642)	(1.505)	
	Discount Rate				0	0.5	
					(0.431)		
Linear	Age	-0.005	-0.12	-0.065	-0.019	-0.182	0.134
		(0.140)	(0.111)	(0.094)	(0.128)	(0.112)	(0.093)
	Holiday	0.939	1.078	0.768	0.939	0.948	1.151
		(0.206)	(0.180)	(0.148)	(0.203)	(0.184)	(0.253)
	Named movies	6	7	6	6	6	6
	Forward Looking	No	No	No	Yes	Yes	No
	Observations	1,432	1,611	1,432	1,432	1,432	1,432

Column 1 adds a random coefficient on *Action*. Column 2 uses 7 named movies instead of 6. Column 3 assumes the moment condition on the first difference of ξ_{jt} rather than the level. Column 4 assumes consumers have perfect foresight and estimates the discount rate. Column 5 holds the discount rate constant at 0.5. Column 6 uses importance sampling.

Table A3: Robustness for welfare and market share effects of the import liberalization from 2012

	Quantities					Percentage Increase from Liberalization			
			SARFT Model				SARFT Model		
Baseline	All movies	Bottom 14	Pre-2012	Post-2012	Top 14	Bottom 14	Pre-2012	Post-2012	Top 14
Annual Admissions	638.13	625.82	607.77	605.09	548.58	1.97	5	5.46	16.32
Annual Admissions of Foreign Movies	280.91	256.76	229.47	225.71	148.68	9.4	22.48	24.53	88.94
Annual Admissions of Competing Foreign Movies		246.83	211.41	206.42	110.51	-3.87	-7.88	-8.56	-25.67
Annual Admissions of Domestic Movies	357.22	369.06	378.3	379.38	399.9	-3.21	-5.57	-5.84	-10.67
Annual Consumer Welfare (Util)						4.03	8.27	8.88	21.16
With replacement									
Annual Admissions	638.13	626.77	609.32	606.68	551.43	1.81	4.73	5.19	15.72
Annual Admissions of Foreign Movies	280.91	257.87	231.29	227.54	151.87	8.94	21.51	23.54	84.97
Annual Admissions of Competing Foreign Movies		246.83	211.41	206.42	110.51	-4.28	-8.61	-9.3	-27.23
Annual Admissions of Domestic Movies						-3.17	-5.5	-5.78	-10.6
Annual Consumer Welfare (Util)	357.22	368.9	378.03	379.14	399.56	3.77	7.82	8.42	20.36

The left panel presents the outcome in levels in millions, averaged across the three years from 2012-2014. The column *All movies* is the observed outcome when 34 movies 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 9 as described in Section 9.1. 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. *Annual Admissions of Competing Foreign Movies* is the quantity sold for the remaining movies when all movies are available (i.e, under liberalization). The top panel (*Baseline*) uses the model with consumption durability but not recency based on Column 4 of Table 5. The bottom panel (*With Replacement*) uses the same demand model (i.e., without recency) but replaces dropped movies so consumers always have 6 choices. See Appendix D. We do not report standard errors for readability sake but they are below 11 for all items in the left panel and below 0.05 for all items in the right panel.

Appendix B Explaining movie fixed effects

We regress the movie-specific effects from the demand estimation on time-invariant movie characteristics and report the results in Table B1. For this regression, we use only movies that appear in more than one week, so there are 308 observations. In order to account for estimation error in the movie fixed effect, we construct standard errors via a mix of a parametric and non-parametric bootstrap. Specifically, we draw a set of dummy variable coefficients from a multivariate normal with means and covariance matrix taken for estimation of our demand model. This is the parametric step. In this way, we construct 2,000 samples. For each sample, we draw one new dataset of dummies from the sample with replacement. This is the non-parametric step. In this way, we have 2,000 samples that reflect both first and second-stage estimation error. We match dummy coefficients to movie characteristics and perform linear regression in each sample. We report the standard deviation of the coefficient estimates across samples.

Focusing on Column 6, our preferred specification, we see that the number of weeks since the international release 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. Most other coefficients are insignificant, but recall that the heterogeneity in the random coefficient on *enhanced* is large. The coefficient on being foreign is negative, but this 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 5 greatly affects the coefficient on the indicator for being a foreign movie in Table B1.

Table B1: Regression of movie fixed effects on movie characteristics

	1	2	3	4	5	6
Enhanced	0.760	1.009	0.510	0.744	0.730	0.148
	(0.493)	(0.472)	(1.374)	(0.443)	(0.317)	(0.555)
Foreign	9.129	6.945	5.793	8.414	-3.108	-1.793
	(3.355)	(4.780)	(2.899)	(3.932)	(3.381)	(4.664)
Weeks since int'l release	0.002	0.011	-0.012	-0.019	-0.003	-0.027
	(0.011)	(0.016)	(0.015)	(0.012)	(0.010)	(0.014)
Ln(run time)	2.408	3.84	2.888	4.874	2.993	6.022
	(0.648)	(1.058)	(0.729)	(0.929)	(0.805)	(1.165)
Action	0.093	-0.006	0.236	0.425	0.093	0.435
	(0.251)	(0.411)	(0.217)	(0.283)	(0.266)	(0.380)
Comedy	0.339	0.168	0.382	0.524	0.249	0.557
	(0.267)	(0.448)	(0.236)	(0.318)	(0.288)	(0.420)
Drama	-0.131	-0.336	-0.084	0.132	-0.072	0.104
	(0.235)	(0.383)	(0.209)	(0.286)	(0.250)	(0.381)
Constant	30.356	-10.860	17.713	-11.673	-35.065	-65.524
	(14.014)	(20.525)	(13.236)	(14.449)	(16.445)	(18.760)
R^2	0.867	0.671	0.818	0.869	0.379	0.251
Consumption Durability	No	No	Yes	Yes	Yes	Yes
Micro-moments	No	Yes	No	Yes	No	Yes
Recency	No	No	No	No	Yes	Yes

The columns are defined analogously to Table 5. Estimation accounts for estimation error of the dependent variable through the bootstrap. 308 observations.

Appendix C Welfare computation

In this section, we describe how we calculate welfare gains from the increase in choices. First, we describe how we obtain a price coefficient in order to convert utils into a dollar-valued number, and then we describe our welfare calculation.

We construct a price coefficient as follows. First, we choose an elasticity to target, χ^*. In our case, $\chi^* = 1.74$, as described in Section 9. Denote the movie fixed effect of movie j as $\bar{\delta}_j$ and the vector of movie fixed effects as $\bar{\boldsymbol{\delta}}$. The change in utility for movie j resulting from a price change of Δp is $\psi \Delta p$, where ψ is the price coefficient. We compute the elasticity by considering the change in market share from the utility of each product going from $\bar{\delta}_j$ to $\bar{\delta}_j + \psi \Delta p$. We calculate the own-price within-period elasticity from our model at each movie in each period and average, denoting the result as $\chi(\bar{\boldsymbol{\delta}}, \Delta P, \psi)$. We find ψ such that $\chi(\bar{\boldsymbol{\delta}}, \Delta P, \psi) = \chi^*$. We set $\Delta P = 10$, about 30% of the average price of a movie in yuan.

For welfare changes, 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 subsection. 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 $\widetilde{C}_{gt}$, $g = 1, .., G_t$. Some choice sets $\widetilde{C}_{gt}$ 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, $\widetilde{C}_{gt}$ 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 effect of import liberalization on consumer welfare, we compute the welfare to consumers with and without the excluded movies, as follows:

$$
\begin{aligned}
CS_i &= \sum_t \sum_{g \in G_t} s_{igt} \ln \left(1 + \sum_{j \in C_{gt}} e^{\delta_{jt} + \mu_{ijt}} \right). \\
\widetilde{CS}_i &= \sum_t \sum_{g \in G_t} \widetilde{s}_{igt} \ln \left(1 + \sum_{j \in \widetilde{C}_{gt}} e^{\delta_{jt} + \mu_{ijt}} \right). \\
\Delta CV &= \int \frac{CS_i - \widetilde{CS}_i}{\psi} dF_i,
\end{aligned}
$$

where CV stands for compensating variation. Note that we have ignored the integrating constant in our definitions of CS_i and $\widetilde{CS}_i$ but because we focus on the difference between them in ΔCV, it would drop out anyway. In order to normalize compensating variation, we compare it to the average annual industry revenue for 2012-2014, denoted IR. Then, the percentage change in consumer welfare is:

$$
\% \Delta CV = \frac{\Delta CV}{IR}.
$$

Appendix D Counterfactual calculations with replacement

In this section, we describe our method for calculating the case in which we ensure 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 10, excluded movies are removed from the choice set, so consumers may have less than six named movies they can choose from. That is, $\boldsymbol{C}_t$ is reduced to five, four, or (for two weeks in our data) even three choices. In the computation presented in Table A3, we replace the excluded movies so that $\boldsymbol{C}_t$ 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 $\boldsymbol{C}_t$ 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 $\boldsymbol{C}_t$ 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 B1), 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 5), 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 a 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 5). We calculate the average difference between the mean utility of the sixth and seventh most popular movies. In particular, let

$\tilde{\delta}_t^{(k)}$ be the mean utility of kth 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 $\min\{\delta_t^{(8)}, \delta^{(6)} - 2\Delta\delta\}$ 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 18). 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 calculations follow exactly as in the baseline case. Note that for all panels of Table 10 and Table A3, we hold α (the calibrated price coefficient) at the same value rather than recalibrating it for each specification. We do this so we can focus on change due to the model rather than the price coefficient.