Something to Talk About: Social Spillovers in Movie Consumption

Duncan Sheppard Gilchrist

Wealthfront

Emily Glassberg Sands

Coursera

We exploit the randomness of weather and the relationship between weather and moviegoing to quantify social spillovers in movie consumption. Instrumenting for early viewership with plausibly exogenous weather shocks captured in LASSO-chosen instruments, we find that shocks to opening weekend viewership are doubled over the following five weekends. Our estimated momentum arises almost exclusively at the local level, and we find no evidence that it varies with either ex post movie quality or the precision of ex ante information about movie quality, suggesting that the observed momentum is driven in part by a preference for shared experience, and not only by social learning.

I. Introduction

Economists have long recognized that an individual's demand for a good may depend on the consumption of the good by others, either because

We would like to thank Natalie Bau, Victor Chernozhukov, David Cutler, Mark Duggan, Liran Einav, Hank Farber, Peter Ganong, Ed Glaeser, Claudia Goldin, Christian Hansen, Larry Katz, Mike Luca, Jeffrey Miron, Markus Mobius, Ariel Pakes, Amanda Pallais, Andrei Shleifer, Fanyin Zheng, and seminar participants at Harvard's labor and public economics lunch for their helpful comments and suggestions. We additionally thank Adam Presser for sharing his insight into the movie industry and Kevin Garewal for his assistance in procuring the data. We are grateful to Jesse Shapiro and five referees for their thoughtful comments on earlier drafts. The title is inspired by Bonnie Raitt's 1991 hit song. Data are provided as supplementary material online.

Electronically published August 30, 2016
[*Journal of Political Economy*, 2016, vol. 124, no. 5]

of direct externalities in utility (Becker 1991) or because of learning from the decisions or experiences of others (e.g., Bikhchandani, Hirshleifer, and Welch 1992; Young 2009).[1] However, despite a long tradition of theoretical research, because of the inherent difficulty in separating causal from noncausal stories for correlation of behavior in social groups (Manski 1993), there remains little clean evidence of social spillovers in consumer demand and even less evidence that separates different economic explanations.[2]

We explore social spillovers in the consumption of a major entertainment good, in-theater movies. The thought experiment is simple: with all other characteristics of a movie held fixed, does an individual's demand for the movie depend on whether others have seen it? We begin by presenting a simple theoretical model of moviegoing in the presence of network externalities, that is, a preference for shared experiences. In the model, the utility of movie attendance at any point depends both on movie quality and on the movie's prior viewership. Sales in one weekend increase the utility of attendance, and thus viewership, in subsequent weekends. The model also shows that momentum generated by network externalities can be independent of other movie characteristics, including both movie quality and the precision of prior information about quality.

To test for and quantify network externalities in movie consumption, our empirical strategy is to exploit the randomness of weather and the relationship between weather and moviegoing.[3] In brief, we instrument for opening weekend viewership with unanticipated and plausibly exogenous weather shocks that weekend. We then estimate the effect of exogenous shocks to opening weekend viewership on viewership in later weekends. Our results show that a shock to opening weekend viewership is doubled over the following five weekends.

By exploiting weather shocks in a movie's opening weekend as a plausibly exogenous source of variation in opening weekend viewership, our empirical specifications avoid many possible confounds.[4] In the first stage,

[1] Other prominent works in the learning literature include Banerjee (1992), Ellison and Fudenberg (1995), McFadden and Train (1996), and Bikhchandani, Hirshleifer, and Welch (1998).

[2] Notable exceptions include a range of observational and experimental studies that have found evidence of information stories driving convergent, or herd, behavior (e.g., Scharfstein and Stein 1990; Welch 1992; Çelen and Kariv 2004; Munshi and Myaux 2006). A large literature in industrial organization also studies the effects of network externalities in platform markets, particularly with regard to the market power they can induce through consumer lock-in (see Farrell and Klemperer [2007] for an overview); and at the firm level, Katz and Shapiro (1986) analyze technology adoption in the presence of network externalities.

[3] Dahl and DellaVigna (2009) document that weather shifts movie sales.

[4] In much of the existing literature analyzing motion picture demand, researchers deal with potentially confounding unobservables by conditioning on opening weekend audience size and then explore how things like reviews or awards shift the demand curve in later weeks (see, e.g., Motiere and Mulligan 1994; Prag and Casavant 1994; Eliashberg and Sawhney

we instrument for opening weekend viewership with weather shocks occurring that weekend. Controlling for general seasonality, these unanticipated weather shocks are orthogonal to unobserved demand and supply shocks and to movie quality. In the second stage, we estimate the effect of instrumented opening weekend viewership on viewership in later weekends. To account for seasonality in movie demand and supply, we define viewership throughout as audience size conditional on year, week of year, day of week, and holiday fixed effects. To account for any autocorrelation in weather, we also condition viewership in weekends subsequent to opening on contemporaneous weather.

Using weather as an instrument is appealing in this setting both because weather is unpredictably variable and because it has real effects on behavior. Instrumenting with weather effectively, however, is nontrivial in part because the set of weather measures is large, particularly at the national level.[5] The risk of either overfitting the first stage (e.g., by including all potential instruments) or data mining (e.g., by hand-picking some instruments and excluding others without objective reason) makes careful variable aggregation and selection methods crucial in this setting. Given the large set of potential weather instruments, we implement least absolute shrinkage and selection operator (LASSO) methods as developed in Belloni, Chernozhukov, and Hansen (2011) to estimate optimal instruments in linear instrumental variable (IV) models with many instruments.

Instrumenting for opening weekend viewership with weather shocks that weekend, we find strong evidence of large and persistent momentum due to network externalities in consumption. For 100 weather-induced additional viewers opening weekend, we observe almost 50 additional viewers in the second weekend and almost 30 in the third. The impact of opening weekend viewership on viewership in subsequent weekends falls approximately exponentially over time, yet the aggregate impact remains large: by the sixth weekend, cumulative momentum has yielded more than one subsequent viewer for each additional viewer during opening weekend. Although our preferred estimates are generated using LASSO-chosen instruments, instrumenting with intuitive, hand-selected instruments yields similar results. We additionally present evidence that our instrument is indeed exogenous and that the observed shift is a demand-side phenomenon: prices are fixed and supply responses (e.g., adjustments to the number

1996; Donihue et al. 2001; Moul 2007). In an insightful twist, Moretti (2011) uses the number of opening theaters as a proxy for expected demand and shows differential momentum from positive and negative shocks to moviegoing as evidence of social learning about movie quality. Such approaches cannot, however, speak to network externalities.

[5] Consider a simple Google search of "02138 weather," which yields a wealth of information including Cambridge, MA's hourly maximum temperature, probability of precipitation, humidity, wind speed, and cloud cover. Moreover, such weather measures are available for each of thousands of weather stations.

of screens on which the movie shows or changes in its duration in theaters) can explain little, if any, of our estimated momentum.

Do these network externalities exist predominantly at the local level (e.g., through conversations among friends) or at the national level (e.g., by way of national media coverage of box office sales)? We proxy for local viewership with metropolitan statistical area (MSA) level Google search volume data and show that the lion's share of our estimated momentum is bred within an MSA.

We also test for evidence of an important alternative explanation for the momentum we observe: learning. The theoretical literature on learning models is large and varied in both its assumptions and its predictions. To make progress, we focus on two intuitive predictions of prominent learning models: (1) in the presence of social learning, shocks to viewership should induce stronger momentum for high-quality movies than for low-quality movies; and (2) in the presence of observational learning, shocks to viewership should induce stronger momentum for movies with more diffuse ex ante priors. Importantly, these predictions need not arise from a model of network externalities and indeed are not implied by our network externalities model. We proxy for realized vertical quality with critic reviews and for the precision of information about quality with production budgets, and we find no evidence that our estimated momentum varies along either dimension. Although our estimates do not rule out some role for learning, taken together the results suggest that the observed momentum is driven in part by a preference for shared experience, and not only by learning.

Finally, we analyze the economic implications of our results. We first ask from where the marginal viewers induced by network externalities are coming. Looking across all movies available in theaters simultaneously, we show that these viewers are predominantly substituting across movies rather than across activities. Second, we leverage our framework to ask whether certain groups are disproportionately affected. We find that our estimated momentum is nearly 50 percent stronger among viewers of child-friendly movies (Motion Picture Association of America [MPAA] rating of G or PG) than among viewers of adult-oriented movies (PG-13 or R). Third, we consider the aggregate magnitudes and show that opening weekend weather shocks explain 9 percent of the variation in opening weekend ticket sales and, because of social spillovers, a full 11 percent of the variation in total ticket sales over the course of a movie's 6-week run in theaters. This is sizable relative to other demand shifters and represents an estimated one-fifth of the effect of advertising.

While theorists have long posited that network externalities may contribute to convergent choices across individuals, empirical evidence on network externalities is limited. Existing work in nonexperimental settings focuses predominantly on the role of learning: Moretti (2011) analyzes social learning in moviegoing, and Cabral and Natividad (2016) provide evidence that top-selling movies in opening weekends tend to earn significantly more in subsequent weekends primarily because of increased awareness.[6] Analyzing social learning in other media, Chen (2008) finds evidence of herd behavior in online book purchasing, and Sorensen (2007) similarly tells an information story in the book market, identifying off of accidental placement on the *New York Times* best seller list.

The work perhaps closest to our own is an experiment by Bursztyn et al. (2014), which finds that investors' stock choices are influenced by network externalities: randomizing both whether an investor's planned stock purchase decision is shared with another investor and whether that purchase is executed, they show that investors value actually owning the same stocks as others (over and above the information contained in others' planned stock purchases). We take a similar approach but identify off of quasi-random variation in early adoption driven by exogenous weather shocks, leveraging LASSO methods to facilitate optimal instrument selection. While existing research using LASSO is predominantly theory and simulation based, Chen and Yeh (2012) and Chen and Sethi (2012) similarly implement LASSO techniques for IV selection in applied settings.

The remainder of the paper proceeds as follows: We open in Section II with a simple model of moviegoing in the presence of network externalities. We then describe our movie and weather data (Sec. III), followed by our empirical approach and first-stage results (Sec. IV). In Section V, we present our base case estimates of network externalities in moviegoing, along with a series of robustness checks; proxying for local moviegoing using MSA-level Google search data, we then show that the observed momentum is bred principally at the local level. Section VI considers learning as an alternative explanation for our results, Section VII explores the economic implications, and Section VIII presents conclusions.

II. A Simple Model of Network Externalities in Moviegoing

In this section, we present a simple model designed to illustrate the mechanism through which network externalities can shape demand for movie tickets. The model yields three predictions, which we return to in subsequent empirical sections. Our framework is highly stylized and does not capture all features of the market; we discuss its limitations below.

[6] Moretti's (2011) empirical analysis actually tries to rule out network externalities in moviegoing; however, as we show in online app. sec. F, his results are not robust to the inclusion of seasonal controls.

A. *Preliminaries*

There is a unit mass of consumers who value both seeing a good movie and sharing in the experience of moviegoing. In particular, the utility of attending a given movie is increasing both in the movie's quality and in its cumulative local viewership. For simplicity, we focus here on demand in a single locality.

The setup is as follows: Before a movie's release, each consumer i observes its quality, α, and her idiosyncratic valuation of viewing, e_i. The consumer attends the movies on her own, but each weekend she learns how many others in her locality attended in prior weekends. Demand for tickets is increasing in the movie's cumulative prior local viewership, denoted CPV_w. (Cumulative prior viewership on the opening weekend, CPV_1, is zero by definition.)

In any weekend a movie is showing, the consumer compares her utility of attending that movie with the opportunity cost of viewership, c, and attends if the difference is at least zero. Altogether, consumer i's utility from viewing the movie on weekend w is

$$U_{iw} = \alpha + e_i + \lambda \mathrm{CPV}_w, \tag{1}$$

where the parameter $\lambda \geq 0$ captures the extent to which consumers value shared experiences. That is, λ parameterizes the strength of network externalities: if $\lambda = 0$, then the value of shared experiences drops out of the utility specification completely, while at larger λ's consumers receive additional utility from moviegoing the more others have already attended.

To close the model, we assume that the idiosyncratic valuations of viewing, e_i, are uniformly distributed on $[0, 1]$, that consumers receive positive utility only from their first viewing of a movie, and that consumers are myopic.[7] These assumptions make the analysis tractable, ensure that demand for movies is decreasing in weekends since release, and rule out the possibility that consumers choose to delay viewing until more of their peers have already viewed. Finally, we assume that $\lambda < 1$ and $\alpha - c \in [0, 1]$, so cumulative viewership is always at least zero and never exceeds one.

Before turning to the analysis and predictions, we note that this model is deliberately held simple. Importantly, the model assumes that movie qualities are public and certain. Incorporating uncertainty in movie quality, α, on its own does not change our results as long as we do not also allow consumers to have private information about movie quality. However, that consumers do not have heterogeneous beliefs about movie quality

[7] Whereas existing learning models of innovation adoption have generally assumed either that learning is social but consumers are myopic (e.g., Ellison and Fudenberg 1993; Young 2009) or that consumers are forward looking but that information arrives exogenously (e.g., Jensen 1982), a new literature (e.g., Frick and Ishii 2016) explores forward-looking social learners.

explicitly rules out the potential for consumers to learn from one another (e.g., as in Moretti [2011]). We return to learning models, and the differences in their assumptions and predictions, in Section VI.

B. Analysis

Opening weekend viewership for the movie is simply given by the range of consumers who receive positive utility from attendance without any utility gain from prior viewership:

$$V_1 = 1 - (c - \alpha). \tag{2}$$

Working forward, cumulative prior viewership in subsequent weekends $w > 1$ is then

$$\mathrm{CPV}_w = 1 - (c - \alpha - \lambda \mathrm{CPV}_{w-1}), \tag{3}$$

and solving the recursion yields

$$\mathrm{CPV}_w = [1 - (c - \alpha)] \sum_{\tau=0}^{w-1} \lambda^{\tau}. \tag{4}$$

Thus, viewership in weekend w is given by

$$V_w = \lambda^{w-1}[1 - (c - \alpha)]. \tag{5}$$

This simple model yields three main predictions.

PREDICTION 1. The ratio of viewership in weekend $w > 1$ to viewership opening weekend decreases exponentially in weeks since opening:

$$\frac{V_w}{V_1} = \lambda^{w-1}. \tag{6}$$

Viewership in a given weekend is a constant fraction λ of viewership in the weekend prior. We return to this prediction in Section V.

PREDICTION 2. Opening weekend viewership V_1 is increasing in movie quality α, that is, $\partial V_1 / \partial \alpha > 0$, but the ratio V_w / V_1 between viewing in a subsequent weekend and viewing in the opening weekend does not depend on α, that is, $\partial (V_w / V_1) / \partial \alpha = 0$.

We return to this prediction in Section VI, where we compare the predictions of our model to learning models.

PREDICTION 3. Stronger network externalities increase viewership in subsequent weekends relative to the opening weekend:

$$\frac{\partial}{\partial \lambda}\left(\frac{V_w}{V_1}\right) > 0 \quad \text{for } w > 1. \tag{7}$$

We return to this prediction in Section VII, where we explore heterogeneity in the magnitudes of momentum across audience demographics.

III. Measuring Weather and Moviegoing

A. *National Ticket Sales and Nationally Aggregated Weather Measures*

Our national box office data come from Box Office Mojo, a reporting service owned by Internet Movie Database (IMDb), and include total US daily ticket sales. In the weeks just following release (when a movie can generally be viewed exclusively in theaters), box office data provide an excellent measure of a movie's total sales.[8] Our ticket sales sample comprises all movies wide-released in US theaters between January 1, 2002, and January 1, 2012.[9] We track audience sizes during the 6 weeks following the date of wide release. To avoid truncation issues, the 19 percent of movies that do not last at least 6 weeks in theaters are excluded from our main analysis.[10] We focus throughout on weekend (Friday, Saturday, and Sunday) audiences since these are most responsive to weather shocks and weekend audiences account for the vast majority (over 75 percent) of ticket sales.[11]

Figure 1 shows average daily ticket sales and average daily ticket sales per screen for each of the first six weekends in theaters. Panel A plots averages across the 1,381 movies in our sample. Average daily ticket sales exceed 1 million during opening weekend but fall off exponentially in subsequent weeks. The modal number of new movies per weekend is two, though some weeks have no new releases and others have as many as five. Since our weather instruments are at the daily level, in our analyses

[8] Though a few distributors have tried experimenting with simultaneous release in theaters and on home video, the vast majority do not release on home video until months (usually 3–4) after the end of the theatrical release. Additionally, although we do not observe viewership of pirated movies, as long as an individual's demand for the pirated version does not fall the more others have seen the movie in theaters, then at worst our estimated network externalities would be biased downward.

[9] We follow Corts (2001) and Einav (2007) in defining as "wide released" any movie that ever showed on 600 or more screens, and we omit from the sample the less than 1 percent of movies that never reached wide release. For the 20 percent of movies in our sample that start with a limited release before reaching wide release, we again follow Einav in defining the wide-release date as the first date on which the movie is shown on more than the maximum of 400 screens and 30 percent of the eventual maximal number of screens for that movie. (Excluding the limited-release movies from the sample does not substantively change our results.) Although box office data are available for earlier years, we focus on the post-2001 period because for earlier years most ticket sales data are reported only at the movie by week level.

[10] We return to them when examining supply responses in online app. sec. D and show that our results are robust to their inclusion. Leveraging Box Office Mojo's reports on the number of screens on which the movie shows weekly, in online app. sec. D we also analyze any supply shifts that might affect our observed quantity effects. Combined with total ticket sales quantities, the number of screens (supply) facilitates an isolation of demand shifts since ticket prices are generally fixed.

[11] Restricting to weekend audiences is standard in this literature (see, e.g., Dahl and DellaVigna 2009).

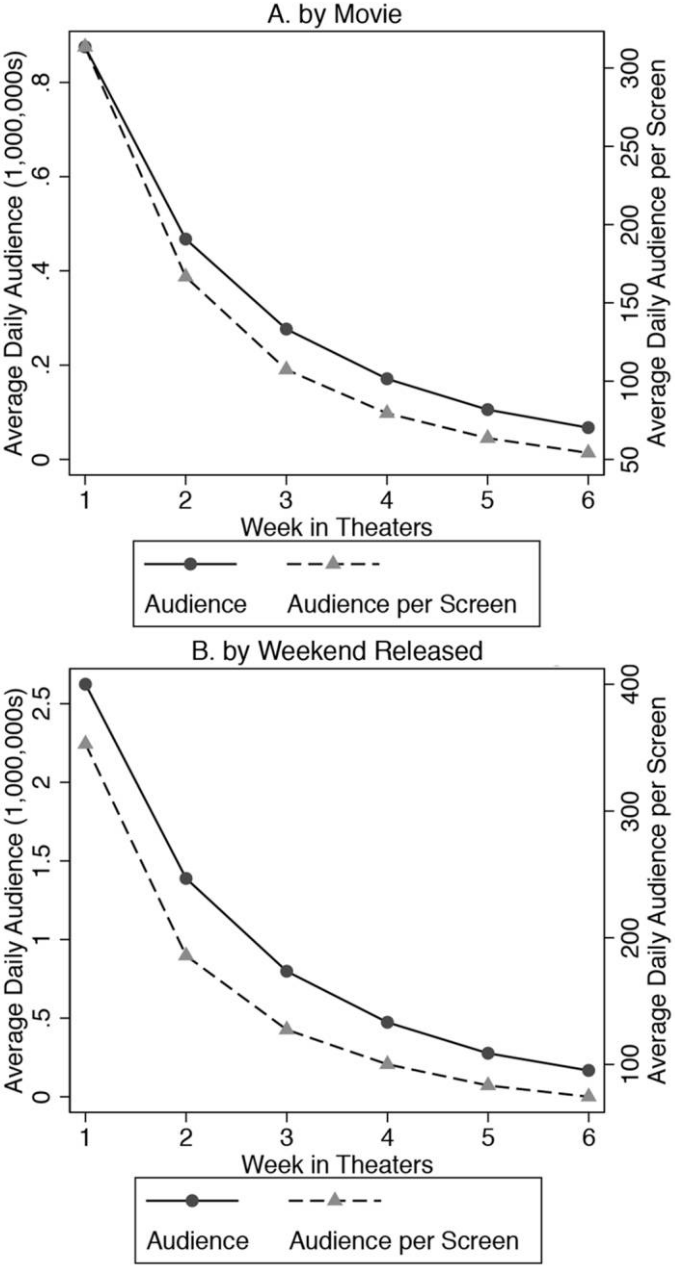

FIG. 1.—Average audience sizes by week in theater. For our sample of 1,381 movies, in panel A we plot average daily ticket sales (in millions) and average daily ticket sales per screen for each of the first 6 weeks in theaters. In panel B, we sum across movies released in the same weekend and report average daily ticket sales and average daily ticket sales per screen for each of the first 6 weeks after release. Here, and throughout our analysis, we restrict to weekend (Friday, Saturday, Sunday) audiences.

we group movies by the weekend on which they were released.[12] Our unit of observation for audiences, then, is at the opening weekend by date level. In our 11-year sample we observe 557 opening weekends, or 1,671 opening weekend days. Panel B of figure 1 plots the average of daily ticket sales (and ticket sales per screen) at the release weekend level. The average audience flocking to new releases is 2.5 million viewers. The corresponding number for movies in their second weekend is just over 1.3 million; this falls to about 200,000 by the sixth weekend in theaters.

Our nationally aggregated weather measures reflect the percentage of movie theaters in the country experiencing a particular type of weather. The raw data are from Weather Underground, a commercial provider of real-time and historical weather information online, with most US data coming from the National Weather Service. We observe daily weather measures for each of 1,941 US weather stations and focus on three weather measures: maximum temperature, precipitation, and the interaction of temperature and precipitation.[13] To reduce the effect of possibly spurious outliers, we winsorize our temperature and average hourly precipitation measures at 1 percent levels. Then, to facilitate national aggregation, we create temperature dummies in 5° Fahrenheit bins and precipitation dummies in 0.25 inch per hour bins, as well as indicators for any snow or any rain.[14] Our nationally aggregated weather measures are simply, for each measure, the weighted average of that measure across weather stations. Weights are assigned to weather stations annually on the basis of the percentage of total movie theater establishments to which the weather station is matched.[15]

[12] Almost all movies are released on Fridays; a few are released on Wednesdays. For Wednesday releases, we omit the first two daily observations, thereby treating the first Friday after opening as the opening date. Grouping by opening weekend is thus equivalent to grouping by opening date or by opening week.

[13] We use maximum temperature (rather than minimum temperature) because we expect much of weather's impact on moviegoing to be driven by its effect on alternative afternoon (heat of the day) activities; evening substitutes for movies such as dinners and indoor parties are not as heavily weather dependent as afternoon barbecues and pool time.

[14] Our motivation for dummying out before aggregating is perhaps best shown by example. Suppose that the population lived in equal numbers in two cities, Los Angeles and Boston. On a particular fall day, Los Angeles had a maximum temperature of 105° while Boston had a maximum of 55°. If we aggregated nationally by simply taking the weighted average across cities, we might erroneously conclude that the country experienced a beautiful (80°) day when in fact half the country was cool and half was hot. Our weather measures are designed to capture this variation in temperature.

[15] From the US Census Bureau's annual Zip Code Business Patterns data, we observe for each year from 2002 to 2011, inclusive, the number of theater establishments in each zip code. Since the 2012 data were not available at the time of writing, we proxy for the 2012 establishment numbers with those from 2011. Though the "movie industry" spans across multiple six-digit North American Industrial Classification System (NAICS) codes, we include only establishments with NAICS code 512131, i.e., motion picture theaters (except drive-ins). We match each zip code (and all of its movie theater establishments) to the weather station that is nearest in great-circle distance to the zip code's center, conditional

B. *Proxying for Local Moviegoing with Google Search Data*

We proxy for local moviegoing using search data from Google Trends, which provides a measure of search activity on Google pertaining to specific keywords and topics.[16] Google search data have been used by researchers to proxy for behaviors ranging from turning up at the polls to maltreating children, and even for sentiments such as racial animus (see, e.g., Stephens-Davidowitz 2013a, 2013b, 2014). While search volumes do not perfectly capture movie consumption, they should provide a strong proxy since many people search online for showtimes before attending a movie. We find national searches to be highly predictive of national sales: the correlation between abnormal national searches and abnormal national moviegoing is .74.

We proxy for local moviegoing with Google Trends search data at the MSA by day by topic level (the most granular level at which it is made publicly available) and employ Google's topic classification engine to classify searches as pertaining to particular movies.[17] Smith, Stephens-Davidowitz, and Varian (2015) similarly use local searches for movies as a proxy for local moviegoing (to study the impacts of Super Bowl advertising on movie consumption). Appendix B details our data collection methods and variable definitions. In brief, our MSA by day by movie level search measure corresponds to the *Z*-score of search volume within that MSA. We focus on the five MSAs with the most complete Google data: New York, Los Angeles, Chicago, San Francisco, and Washington, DC.[18]

C. *Proxying for Movie Quality and the Precision of Ex Ante Information about Quality*

In Section VI, we look for evidence of learning driving our estimated momentum; there, we ask whether the observed momentum is stronger for higher-quality movies or for movies for which there is less precise ex ante information about quality.

We proxy for realized movie quality with ratings by expert reviewers. The ratings come from IMDb's top 1,000 voters, a group characterized by IMDb as "the 1000 people who have voted for the most titles in [their]

---

on that distance being no greater than 160 kilometers. For the years in our sample, less than 1 percent of establishments fall outside a 160-kilometer radius of any weather station.

[16] Google is the dominant player in the US search market, with approximately a 68 percent market share at present.

[17] We assume that weather shocks do not change how search behavior and ticket sales relate. We probe this assumption in online app. fig. B.1, which shows how the residual from a regression of abnormal national ticket sales on abnormal national searches is related to national weather shocks in 5° increments from 50° to 95°. All estimated coefficients are close to zero, and none is statistically significant.

[18] Since Google provides only unitless search figures, we are unable to directly compare search volumes across MSAs and instead rely on MSA-level *Z*-score normalizations.

ratings poll."[19] We cut movies into quality terciles on the basis of expert ratings.[20] This dimension captures an important feature of demand: on average, top tercile movies sell nearly 8 million tickets over the course of the first 6 weeks, while bottom tercile movies, on average, sell just 4.6 million tickets over the same period.

We proxy for the precision of information available ex ante about a movie with the movie's production budget. Advertising budgets are generally set as a fixed percentage of production budgets (Einav 2007), and as we describe below, we find that production budgets are positively correlated with the precision of prior knowledge.[21]

Adapting the methodology of Moretti (2011), we define a movie's "surprise" sales as the total ticket sales in weekends subsequent to opening not predicted by the number of opening weekend theaters, that is, the residual from a regression of the log of total sales in weekends subsequent to opening on the log of opening weekend theaters. This reflects the portion of ticket sales that was not predicted by what was known about a movie prior to its release, where the ex ante prediction is captured by the number of theaters that, operating as rational profit maximizers, chose to screen the movie in its opening weekend.[22]

When we split production budgets by tercile, the standard deviations of surprise for movies in the bottom (below $29 million), middle (between $29 million and $48 million), and top (in excess of $48 million) terciles are 0.92, 0.83, and 0.75, respectively. Panel A of online appendix figure B.2 shows surprise plotted against production budget, and panel B shows separate kernel densities of surprise for movies in the top and bottom terciles of production budget. The distribution for movies in the bottom tercile has less mass in the middle and more in the tails, and a two-sample Kolmogorov-Smirnov test shows that the distributions of surprise for top and bottom terciles are statistically significantly different at the 10 percent level (p-value = .099), even though Kolmogorov-Smirnov tests are known to be insensitive to differences in the tails (Conover 1999). We infer that

[19] This is assuredly not the only dimension of quality that might matter, but it is readily quantified. IMDb notes that they "don't disclose the number of votes required for a person to make this list nor can [they] confirm or deny who is on the list." All movies in our sample have had at least a full year to accrue votes and on average have been rated by 483 of these top 1,000 voters.

[20] Movies with an average top 1,000 voter rating of 6.3 or above fall in the top tercile, while movies with a rating of 5.6 or below fall in the bottom tercile.

[21] From among the 1,381 movies in our main sample, we have production budgets from IMDb for 88 percent.

[22] We exclude opening weekend sales from the specification so as to focus only on those ticket sales that occur after a movie's quality has been publicly realized. Our main results always focus on subsequent weekend sales relative to first-weekend sales; subsequent weekend sales are those that, at least to some degree, incorporate realized quality and not merely expected quality. However, we note that the results are similar if opening weekend sales are included.

production budgets do serve as a meaningful proxy for uncertainty about quality prior to a movie's release.

IV. Empirical Model

A. *Instrumenting for Viewership with Weather Shocks*

Given the indoor nature of moviegoing, it is perhaps not surprising that a day's weather is an excellent predictor of viewership. When it is beautiful out, there are generally fewer moviegoers; when the weather is less ideal, ticket sales tend to be higher. That is not to say, however, that the observed relationship is causal. As Einav (2007) demonstrates, the seasonality of viewership is driven by seasonality both in underlying demand and, since the supply side takes into account expected demand in timing releases, in the number and quality of movies available in theaters.

Because seasonality is an important component of both the demand and supply, throughout we condition viewership on year, week of year, day of week, and holiday fixed effects and refer to the resulting residuals as "abnormal" viewership. We denote the viewership on date t of movies that are in their jth week of showing by v_{tj}. To compute abnormal viewership during opening weekend, we first regress viewership in opening weekend, v_{t1}, on a constant and a vector of indicators for day of week, week of year, year, and holidays, which we denote F_t:[23]

$$v_{t1} = \beta_1 + F_t'\Phi_1 + \varepsilon_{t1}. \tag{8}$$

We call the resulting fitted values $\widehat{v_{t1}}$ and define abnormal viewership opening weekend as the difference between realized and predicted viewership:

$$v\_\text{abn}_{t1} = v_{t1} - \widehat{v_{t1}}. \tag{9}$$

We want to instrument for this abnormal viewership opening weekend with contemporaneous weather shocks. Given the natural (and anticipated) seasonality of weather and our desire to capture the unanticipated component, throughout we condition each of our weather measures on the same fixed effects as above. That is, for each weather measure w_k, $k \in \{1, \ldots, p\}$, we estimate

$$w_{tk} = \delta_k + F_t'\Phi_k + \varepsilon_{tk}, \tag{10}$$

where t again indexes the date, k indexes the particular weather measure, and the fixed effects, F_t, are as defined in equation (8). We call the resulting fitted values $\widehat{w_{tk}}$ and define the weather shock w_shock_{tk} as the difference between the realized and predicted weather measures:

[23] Appendix A presents the full set of holidays.

$$w\_\text{shock}_{tk} = w_{tk} - \widehat{w_{tk}}. \tag{11}$$

With our controls for seasonality and time trends in both weather and viewership, these weather shocks are plausibly orthogonal to movie characteristics as well as to other demand and supply shocks.

Figure 2 previews a simplified version of the relationship between weather shocks and abnormal viewership during opening weekend. Each coefficient is the result of a separate regression of abnormal viewership on contemporaneous weather shocks in 5° bins. For exposition, we focus on the common summer range of 60°–95°. Amid unexpectedly beautiful weather (especially 75°–80°), opening weekend ticket sales are lower than would be predicted by seasonality. In the presence of weather that is unexpectedly a bit too cool or too warm, in contrast, audiences are larger. Panel A shows the estimated magnitudes when weather is measured as the percentage of movie theaters unexpectedly (for the time of year) in the given temperature range and when viewership is measured in residualized ticket sales. Each plotted coefficient, then, represents estimated abnormal viewership when all (vs. no) theaters are unexpectedly in that temperature range. When 10 percent of theaters unexpectedly experience "ideal" temperatures (i.e., in the 75°–80° range), for example, aggregate viewership to movies opening that weekend is about 300,000 per day lower than expected (one-tenth of 3 million). With daily aggregate viewership to opening movies averaging about 2.5 million, this corresponds to more than a 10 percent reduction in viewership.

To facilitate comparison of effect sizes across temperature ranges, panel B shows the corresponding results when each weather shock is normalized to be mean zero with unit variance, and residual ticket sales are similarly normalized. This panel also serves to illustrate the sizable role a day's weather can play in that day's ticket sales. For a one standard deviation increase in the percentage of movie theaters unexpectedly in the 75°–80° range, for example, we observe a 0.15 standard deviation decrease in daily viewership of opening movies. This corresponds to about 230,000 fewer tickets per day to opening movies, or about a 10 percent reduction in average ticket sales per day to opening movies (2.5 million).

Although weather shocks are important predictors of abnormal viewership, the large number of potential weather shock specifications makes variable selection methods appealing. We provide further detail on our methods for instrument selection in the following subsection; for now, let us take as given the machine-chosen instrument set, which we denote W^{LASSO}. To obtain the first stage, we run ordinary least squares (OLS) on the LASSO-selected instrument(s):[24]

[24] This is often referred to as post-LASSO; coefficients estimated by post-LASSO differ from those estimated via LASSO because of LASSO's shrinkage bias.

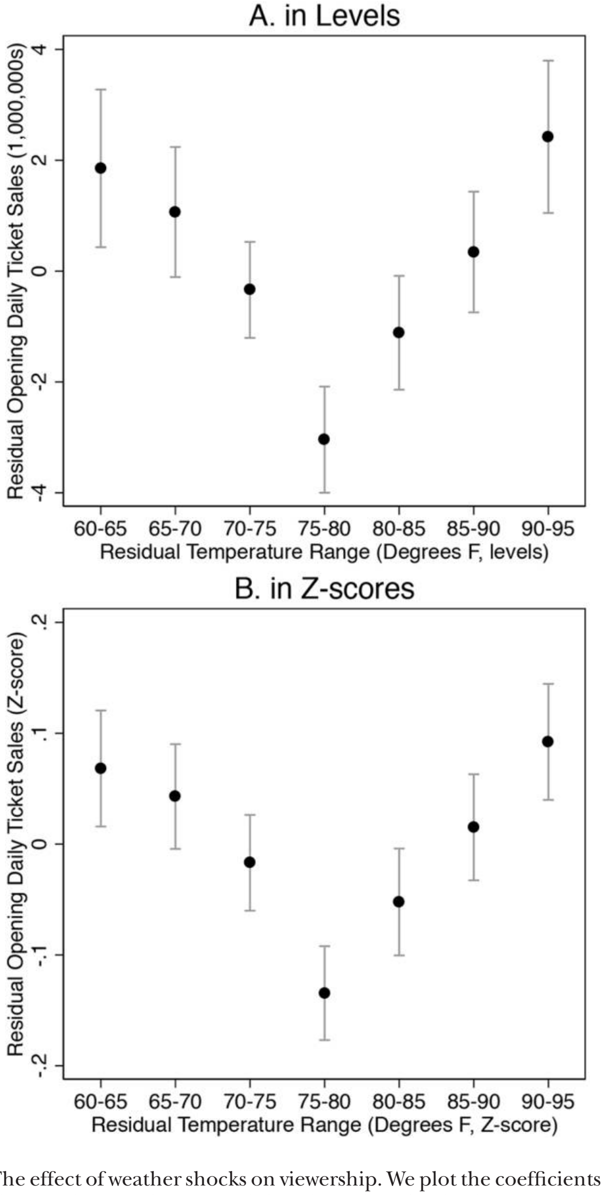

FIG. 2.—The effect of weather shocks on viewership. We plot the coefficients from the regression of abnormal viewership on each listed weather shock, along with the corresponding 95 percent confidence intervals. National weather measures are as described in the text. Panel A shows the relationship in levels: weather shocks are measured as the percentage of theaters unexpectedly in a given temperature bin, and abnormal viewership is measured in number of tickets. Panel B shows the relationship in *Z*-scores: weather shocks are measured as the *Z*-scores of the percentage of theaters unexpectedly in a given temperature bin, and abnormal viewership is measured as the *Z*-scores of the number of (residual) tickets sold. Each plotted coefficient is from a separate regression. Observations are at the date level (1,671 observations).

$$v\_\mathrm{abn}_{t1} = \eta + [W_t^{\mathrm{LASSO}}]'\Omega + \varepsilon_{t1}. \tag{12}$$

We call the resulting fitted values $\widehat{v\_\mathrm{abn}}_{t1}$.

In the second stage, we estimate the relationship between this weather-induced abnormal viewership opening weekend and abnormal viewership in subsequent weekends. We define abnormal viewership in subsequent weekends as viewership conditional on year, week of year, day of week, and holiday fixed effects; given the potential for autocorrelation in weather shocks, we condition also on contemporaneous weather. That is, separately for each week $j > 1$, we first regress viewership on the set of fixed effects and contemporaneous weather:

$$v_{tj} = \beta_j + F_t'\Phi_j + X_t'\Gamma_j + \varepsilon_{tj}. \tag{13}$$

The fixed effects in F_t are as defined in equation (8), and X_t denotes the vector of contemporaneous (date t) weather.[25] We call the resulting fitted values $\widehat{v_{tj}}$ and define abnormal viewership in subsequent weekends as the difference between realized and predicted:

$$v\_\mathrm{abn}_{tj} = v_{tj} - \widehat{v_{tj}}. \tag{14}$$

Finally, to estimate the impact of abnormal viewership opening weekend on abnormal viewership j weeks after opening, we run the second stage separately for each $j > 1$:

$$v\_\mathrm{abn}_{tj} = \mu_t + \theta_j \widehat{v\_\mathrm{abn}}_{t-7(j-1),1} + \varepsilon_{tj}. \tag{15}$$

The estimated momentum in the jth week of showing is $\hat{\theta}_j$.

B. LASSO Instrument Selection and First-Stage Results

Instrumenting with weather is nontrivial in part because the set of potential weather measures is large. Given the issues with either hand-picking a small number of instruments or naively including a large number of instruments, we implement variable selection methods. In particular, we follow Belloni et al. (2011) and implement LASSO methods to estimate optimal instruments in linear IV models with many instruments. The LASSO procedure provides a principled method for instrument selection; in simulation experiments, it performs well relative to recently advocated many-instrument robustness procedures (see Belloni et al. 2012). We follow Belloni, Chernozhukov, and Hansen (2014) in using conventional standard errors and also in adding a constraint on the maximum number of instru-

[25] The vector X_t includes maximum temperature in 10° increments as well as rain, snow, and average precipitation in quarter inches per hour.

ments chosen.[26] Here, we present the machine-chosen instrument sets and the corresponding first stages, and we refer the interested reader to online appendix section C for a brief overview of our LASSO method, which draws heavily on Chernozhukov and Hansen (2013).

We provide LASSO a set of 52 potential weather instruments and denote the final output of the LASSO methodology (i.e., the machine-chosen instrument set) by W^{LASSO}.[27] With a single-instrument constraint, the LASSO-chosen instrument is the 75°–80° measure. Figure 3 shows a histogram of this weather measure; the mass is fairly tightly distributed between −10 percent and +10 percent. Figure 4 shows the corresponding first-stage relationship in a binned scatter plot: the more theaters that are unexpectedly in the 75°–80° range, the lower abnormal viewership is.[28] The first row of panel A of table 1 shows the corresponding first-stage results. Here, 10 percentage points more theaters unexpectedly in the 75°–80° range corresponds to about 300,000 lower daily viewership opening weekend (over 10 percent of average daily viewership for new releases).

For robustness, the remainder of panel A of table 1 shows the first-stage results when we instead constrain to a maximum of two or three instruments or when we constrain to a maximum of one instrument from among a choice set of 10° temperature bins; panel B shows the first-stage results with hand-selected instrument sets. One hand-selected instrument we use is motivated by the intuitive observation (visualized in fig. 2) that the effect of weather on moviegoing is roughly quadratic with a minimum in the 75°–80° range: the instrument is simply the squared difference in temperature from 75° (with the difference set to zero at 55° and 95° to avoid assigning excess weight to outliers).[29] The second set of hand-selected

[26] Conventional standard errors are appropriate as long as the number of selected instruments is not close to the sample size (see Belloni, Chernozhukov, and Hansen [2014] and Cattaneo, Jansson, and Newey [2015] for more detail). We probe the instrument constraint specification choice below and show that our results are not sensitive to different instrument counts.

[27] The full set of potential instruments we provide to LASSO includes the temperature variables in 5° bins for both Saturday and Sunday (18 for each day), indicators for snow and rain (two for each day), and average precipitation in quarter inch per hour variables (six for each day). These weather measures are as defined in Sec. III.A. In our baseline specifications, LASSO always chooses Sunday temperatures, which is consistent with a high volume of daytime (weather-dependent) moviegoing on Sundays. In these specifications, the snow, rain, or average precipitation variables are not chosen.

[28] While the procedure we use assumes a homoscedastic error structure, in a recent paper, Belloni et al. (2014) develop an extension of the LASSO-IV framework that adapts the penalty parameter for the case of a clustered error structure with linear individual-specific fixed effects. We implement their procedure in our setting and find that whether we cluster either by date or by opening weekend (as we do in computing our second-stage standard errors), their new methods consistently yield the same 75°–80° measure as the instrument of choice.

[29] As with all our national instruments, we aggregate across weather stations using the same establishment weighting discussed earlier and similarly condition on day of week, week of year, year, and holiday fixed effects.

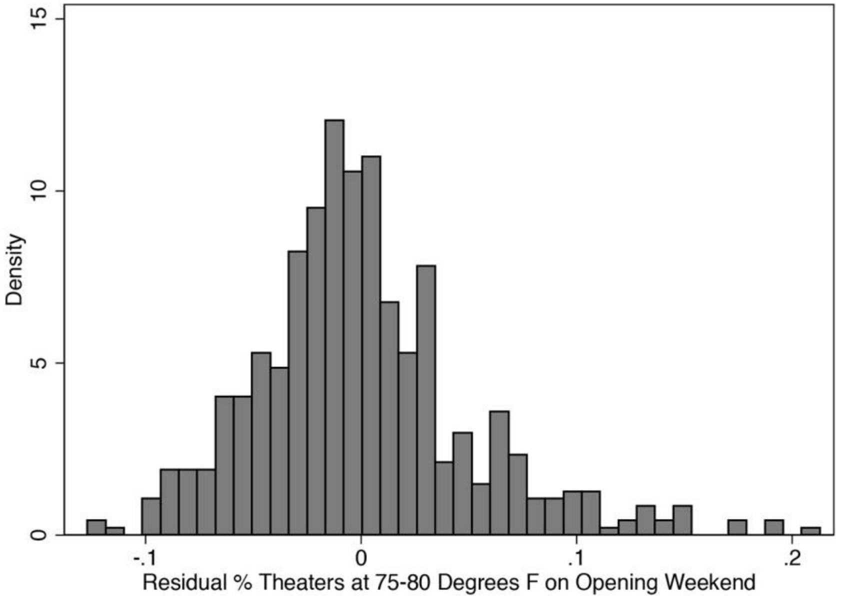

FIG. 3.—Histogram of the instrument. We plot a histogram of the abnormal percentage of movie theaters with weather in the 75°–80° range.

instruments is just the set of all potential weather instruments we supply to LASSO. By any conventional measure of instrument strength, the instruments remain strong in each of these cases, and as we show in the next section, our second-stage results importantly remain largely unchanged across these instrument sets.

V. Momentum from Exogenous Viewership Shocks

A. Base Case Results

The binned scatter plots in figure 5 show the reduced-form relationship between opening weekend weather and ticket sales in subsequent weekends: when more theaters are unexpectedly in the 75°–80° range during a movie's opening weekend, abnormal viewership of that movie is lower not only in that weekend (fig. 4) but also in each subsequent weekend (fig. 5). Implementing the second stage (eq. [15]), we find similarly substantial momentum from exogenous shocks to opening weekend viewership. The first row of table 2 presents our base case IV estimates. The first five columns report the relationship between abnormal viewership opening weekend and subsequent abnormal viewership, separately for each of weekends 2–6. The final column reports the corresponding aggregate relationship, where the outcome is summed across those weekends.

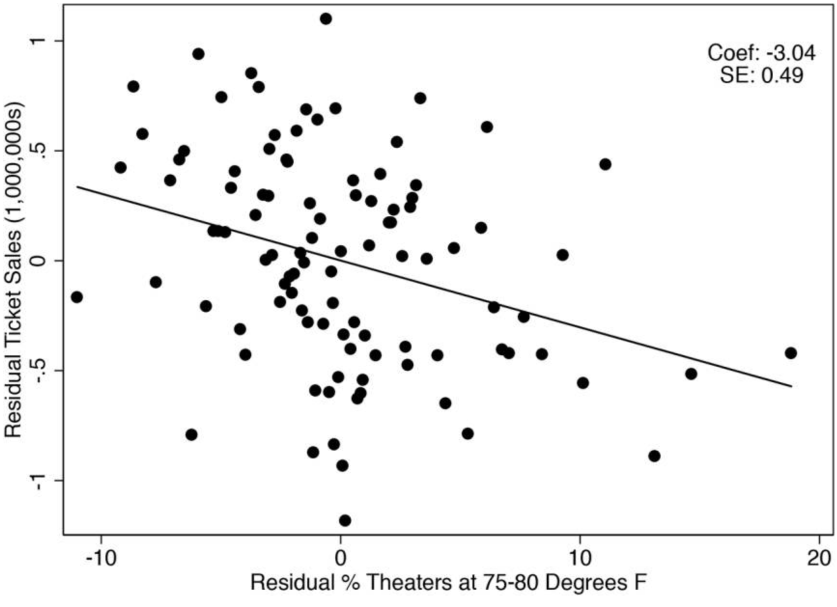

FIG. 4.—First-stage bin scatter. We plot the abnormal percentage of movie theaters with weather in the 75°–80° range against abnormal viewership. For exposition, the weather shock measure is grouped into 100 equal-sized bins; each point corresponds to the mean weather shock and abnormal viewership within a bin. The slope of the line of best fit and the corresponding robust standard error clustered by date are included on the figure.

In the first row, we instrument for abnormal viewership during the opening weekend with contemporaneous weather shocks. One hundred additional viewers on opening weekend yields an estimated 114 additional viewers across the following five weekends. The observed momentum is largest in the weekend immediately following the opening weekend, when nearly half of the total effect is realized; an additional quarter is realized in the third weekend. Although the magnitude of the effect falls off in subsequent weeks, it remains statistically significantly above zero through each of the five subsequent weekends.

The pattern presented in these estimates is consistent with prediction 1 of the model in Section II: viewership in any weekend $w > 1$ is simply a fraction λ^{w-1} of viewership opening weekend. Indeed, a value of $\lambda = 0.5$ in the model predicts that each opening weekend viewer corresponds to 0.5 viewer in the second weekend, 0.25 viewer in the third, and 0.125 viewer in the fourth; by the end of the sixth weekend, this cumulatively predicts about one additional viewer for each viewer opening weekend. This pattern is highly similar to that reported in table 2.

The corresponding OLS estimates, presented in panel B of table 2, are (statistically insignificantly) below the IV estimates. This runs counter to

TABLE 1
LASSO-Chosen and Hand-Selected First Stages

Set of Potential Instruments	Count Constraint	LASSO-Chosen Instrument(s)	Coefficient	*F*-Statistic
	A. LASSO-Chosen Instruments			
5° temperature increments	Choose 1	75°–80°	−3.041*** (.488)	38.80
	Choose 2	75°–80°	−2.635*** (.487)	25.86
		50°–55°	3.419*** (.811)	
	Choose 3	75°–80°	−2.686*** (.488)	20.95
		50°–55°	3.165*** (.837)	
		10°–15°	−2.756** (1.097)	
10° temperature increments	Choose 1	70°–80°	−1.253*** (.319)	15.47
	B. Hand-Selected Instruments			
(Temp − 75°)^2 × [abs(temp − 75°) ≤ 20]			.00449*** (.000824)	29.74
All instruments provided to LASSO in base case			. . .	3.804

Note.—Panel A presents first-stage results for a variety of LASSO specifications. In the first three, the instrument choice set is as follows: national aggregates of maximum temperature indicators in 5° increments (on the interval [10°, 100°]), indicator for snow, indicator for rain, and precipitation indicators in quarter inch per hour increments (on the interval [0, 1.5]). All potential instruments are conditional on the full set of seasonal controls described in the text. From this set, LASSO is set to choose a maximum of one, two, or three instruments, respectively. In the fourth specification, a single instrument is again chosen, but the instrument choice is altered to include the analogous temperature measures in 10° increments instead of in 5° increments. Panel B presents first-stage results for two sets of hand-selected instruments. In the first, the instrument is specified to be the abnormal squared difference between temperature and 75° multiplied by a dummy to ensure that temperatures are within 20° of 75°. In the second, all of the potential instruments that are provided to LASSO in the first three specifications of panel A are included: national aggregates of maximum temperature indicators in 5° increments (on the interval [10°, 100°]), indicator for snow, indicator for rain, and precipitation indicators in quarter inch per hour increments (on the interval [0, 1.5]), again conditional on the full set of seasonal controls described in the text. Observations are at the opening weekend by date level (1,671 observations). Standard errors, clustered at the date level, are in parentheses.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

the intuition that heterogeneity in movie quality could lead to upward-biased OLS estimates of momentum. Since movie quality is certainly heterogeneous (e.g., movie budgets explain 43 percent of the variation in opening weekend viewership), we suggest two possible interpretations. The first is sampling variability: our confidence intervals admit OLS coefficients larger than the IV coefficients (e.g., up to 4 percent larger in the

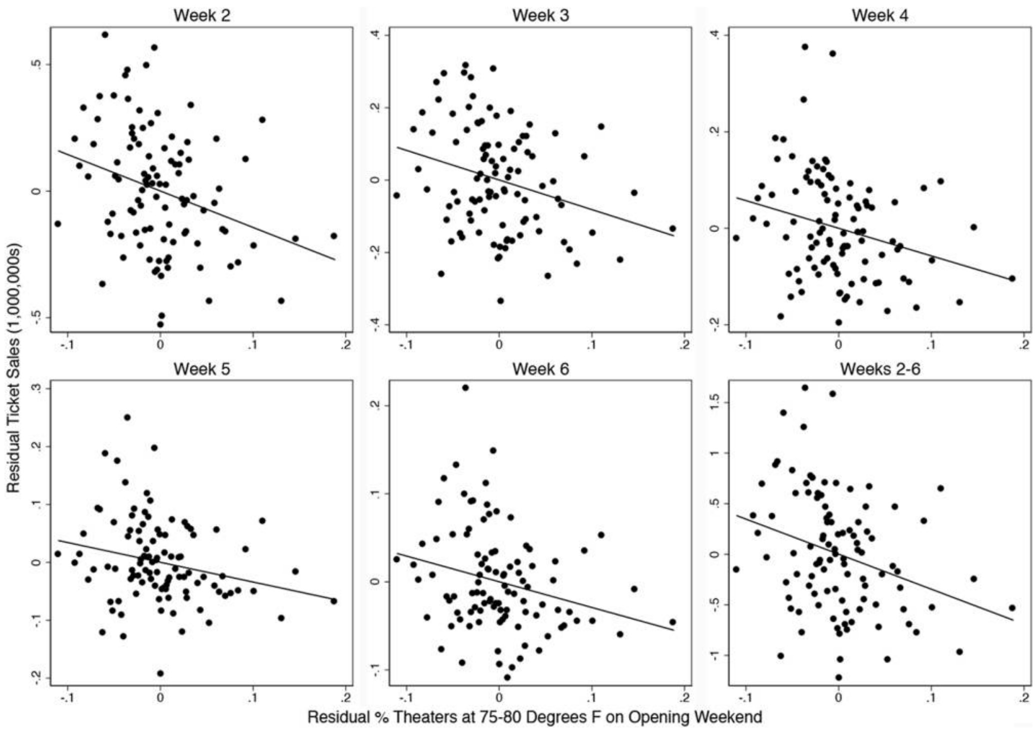

FIG. 5.—Reduced-form bin scatters. We plot the abnormal percentage of movie theaters with weather in the 75°–80° range against abnormal viewership in each subsequent weekend 2–6 and overall across those weekends. For exposition, the weather shock measure is grouped into 100 equal-sized bins; each point corresponds to the mean weather shock and abnormal viewership within a bin.

TABLE 2
Momentum from Viewership Shocks

Instrument	Week 2	Week 3	Week 4	Week 5	Week 6	Weeks 2–6
	A. Instrumental Variables					
LASSO-chosen instrument	.474***	.269***	.188***	.112***	.0960***	1.139***
	(.0474)	(.0360)	(.0287)	(.0203)	(.0162)	(.131)
$(\text{Temp} - 75)^2 \times [\text{abs}(\text{temp} - 75°) \leq 20]$	.547***	.308***	.126***	.0794***	.0402**	1.101***
	(.0509)	(.0401)	(.0290)	(.0204)	(.0171)	(.128)
All instruments provided to LASSO	.475***	.269***	.164***	.121***	.0932***	1.122***
	(.0242)	(.0223)	(.0167)	(.0132)	(.0103)	(.0739)
	B. Ordinary Least Squares					
NA	.423***	.235***	.140***	.0888***	.0630***	.950***
	(.0152)	(.0111)	(.00721)	(.00498)	(.00362)	(.0388)
R^2	.653	.498	.357	.301	.264	.570

Note.—Panel A reports the results of IV regressions of daily abnormal audiences in each later weekend on daily abnormal audiences opening weekend. In the first specification, national weather shock instruments are chosen using the LASSO approach described in the text. The first-stage results are included in the first row of table 1. In the next two specifications, the instruments are not machine chosen: the second specification in panel A presents the results where the instrument is specified to be the abnormal squared difference between temperature and 75° multiplied by a dummy to ensure that temperatures are within 20° of 75°; the third specification presents results obtained by including in the first stage all of the potential instruments that are normally provided to LASSO. The first-stage results for these two specifications are included in panel B of table 1. Panel B reports the corresponding OLS results. Observations are at the opening weekend by date level (1,671 observations). Standard errors, clustered at the date level, are in parentheses.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

model analyzing viewership across all weekends).[30] The second is heterogeneity in momentum effects: whereas the OLS estimates capture the average momentum effect across all abnormal viewers, the IV estimates pertain to the abnormal viewers whose choice was driven by a weather shock. These marginal viewers may be more likely to have friends who are also undecided moviegoers, suggesting that network externalities from their viewership could be stronger than from the average viewership.

B. *Instruments, Clustering Level, and Other Robustness Checks*

We try a number of variants of our base case specification and find that the results are generally unchanged. The second and third rows of panel A

[30] We compute this statistic by estimating the OLS and IV models together in a generalized method of moments framework and report the upper bound of the 95 percent confidence interval of the difference of the OLS momentum estimate less the IV momentum estimate. We normalize this bound into percentage terms by dividing by the IV estimate.

of table 2 show the corresponding results when instrumenting with the two sets of hand-selected instruments discussed in Section IV.B (the squared difference in temperature from 75° degrees and the set of all potential instruments we provide to LASSO), and panel A of online appendix table A.2 shows the corresponding results when LASSO is instructed to choose two instruments or to choose three instruments, or when the potential instrument set is altered to include temperature variables in broader, 10° increments. For parsimony and simplicity, in the remainder of the paper we report results using only the first LASSO-selected instruments but note that the results are essentially unchanged with these alternative instrument sets.

In online appendix table A.1, we show that our results are robust to clustering at the weekend level (panel A), to a more coarse, opening weekend by weekend, unit of observation (panel B), and to the inclusion of second-stage contemporaneous weather controls in the first stage (panel C).[31] In each robustness check, the estimated coefficients closely resemble our base case results, and estimated momentum from exogenous viewership shocks remains highly statistically significant in each week.

C. *Evidence on Exogeneity*

To test the exogeneity of our weather shocks, we ask whether they are correlated with expected demand. In table 3, we follow Moretti (2011) in proxying for expected demand with the number of screens on which the movie opened and control for this measure of expected demand in our specification from before.[32] As Moretti notes, the number of screens is set by profit-maximizing theater owners who have strong incentives to accurately predict opening weekend demand; it should thus summarize well all the information the market has up to the release date about how well the movie will do. For ease of comparison, we reproduce in the first row the results of our main specification. The second row shows the results when we add a control for the number of screens on which the movie opened. Controlling for expected demand, the estimated momentum falls only slightly (and insignificantly); the average change of the point estimates is on the order of 2 percent, and each week's estimated momentum remains large and highly significant. In the third row, we define the outcome variable as abnormal viewership per opening screen. For comparison with our base case, in the fourth row we standardize the coefficients so that the first weekend's coefficient is one. Our estimates again fall only slightly (in-

[31] For example, when the second-stage outcome variable is ticket sales in week 2, week 2 contemporaneous weather controls (as defined in fn. 25) are included also as controls in the first stage.

[32] First-stage results for this specification, as well as for other specifications that extend our base case, are reported in online app. table A.3.

TABLE 3
Momentum per Opening Screen from Exogenous Viewership Shocks

	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6	Weeks 2–6
Tickets (1)	1.000***	.474***	.269***	.188***	.112***	.0960***	1.139***
	(.000)	(.0474)	(.0360)	(.0287)	(.0203)	(.0162)	(.131)
Tickets, controlling for opening theaters (2)	1.000***	.457***	.257***	.186***	.110***	.0977***	1.107***
	(.000)	(.0516)	(.0398)	(.0320)	(.0227)	(.0182)	(.145)
Tickets per opening theater	1.019***	.353***	.248***	.173***	.120***	.115***	1.010***
	(.240)	(.115)	(.0775)	(.0535)	(.0408)	(.0296)	(.289)
Standardized tickets per opening theater (3)	1.000***	.346***	.243***	.170***	.118***	.113***	.991***
	(.238)	(.114)	(.0767)	(.0529)	(.0404)	(.0293)	(.286)
Differences:							
(1) − (2)	. . .	.017	.012	.002	.002	−.002	.032
		(.070)	(.054)	(.043)	(.030)	(.024)	(.195)
(1) − (3)	. . .	.121	.021	.015	−.008	−.019	.129
		(.123)	(.084)	(.059)	(.045)	(.033)	(.315)

Note.—This table shows results from three different IV specifications. The first row reproduces our base case results from table 2. In the second row, we control for the number of opening theaters, a proxy for expected demand, in both the first and second stages. In the third row, we define the outcome variable as abnormal viewership per opening screen. Throughout, national weather shock instruments are chosen using the LASSO approach described in the text. The first-stage results for the first and third rows are included in the first row of table 1; the first-stage results for the second row are included in the first row of online app. table A.3. For comparison with our base case, in the fourth row we standardize the coefficients so that the first weekend's coefficient is one. Observations are at the opening weekend by date level (1,671 observations). The final two rows report differences in the coefficients across the specifications. Standard errors, clustered at the date level, are in parentheses.

*** Significant at the 1 percent level.

significantly) relative to the base case and remain large and statistically significant throughout. Our second stage indeed appears to be picking up viewership shocks that are orthogonal to expected demand.

We also find no evidence that initial weather-induced shocks to viewership are related to our proxies for movie quality. Online appendix table A.4 reports the number of top 1,000 IMDb voters, and the likelihood of being characterized as high rated and low rated, as a function of (instrumented) opening weekend ticket sales. Weather-induced shocks to viewership do not significantly affect the number of residual votes cast by expert reviewers, nor is the level of expert ratings broadly affected.

Finally, in online appendix section D, we analyze an additional explanation for the observed momentum: supply shifts. For example, it could be the case that opening weekend viewership shocks lead theater owners to increase screenings, which decreases the effective cost of attendance and (endogenously) increases viewership in subsequent weeks. Thus, in that section of the online materials, we present a brief overview of the supply side of the market and show that our estimated momentum is not driven by supplier responses. Given that prices are almost always fixed for in-theater movies, we conclude that demand shifts must be driving the observed quantity effects.

D. *Estimating Momentum Locally*

Our methodology for estimating momentum locally is simply the local analogue of our main methodology outlined in Section V.A. Since our local data vary at the MSA level, we now condition our MSA-level weather and MSA-level search data on MSA-level fixed effects for day of week, week of year, year, and holiday. We use the same LASSO procedure for selecting weather instruments.

One advantage of estimating momentum using local-level search data is that we can more directly control for movie quality and latent demand by controlling for (1) movie fixed effects and (2) local abnormal searches in the weekend before opening weekend. The first of these controls is designed to account for the effects of any national-level characteristics that might affect demand (e.g., movie quality, advertising); the second is designed to account for local demand in advance of movie opening.

The directionality of the local-level first- and second-stage relationships is very similar to that at the national level. Figure 6 (the local analogue of fig. 2) previews a simplified version of the opening weekend relationship between local weather shocks and abnormal local searches (our proxy for abnormal local moviegoing), controlling for abnormal local searches in the weekend prior to opening and for movie fixed effects. Each coefficient comes from a separate regression of abnormal viewership on contemporaneous weather shocks in 5° bins. For exposition, panels A and B show the two largest cities (New York and Los Angeles, respectively), and panel C shows all included MSAs. As at the national level, when the local weather is unexpectedly beautiful, opening weekend ticket sales tend to be lower than would be predicted by seasonality; and when local weather is unexpectedly cool or warm, local audiences tend to be higher.[33]

[33] As a robustness check, we also analyze whether local searches are affected by weather shocks in distant MSAs, defined as MSAs at least 1,000 kilometers away. Since weather shocks at that distance are negatively correlated (e.g., a day that falls between 90° and 95° in New York City is associated with a cooler day in Los Angeles), in addition to the standard controls (abnormal local searches in the weekend prior and movie fixed effects), we also condition on local weather controls (10° temperature dummies, quarter-inch precipitation dummies, and snow and rain dummies, as in our national regressions). The results, presented in online app. fig. B.3, show no notable relationship between local opening weekend viewership and weather shocks occurring in MSAs more than 1,000 kilometers away.

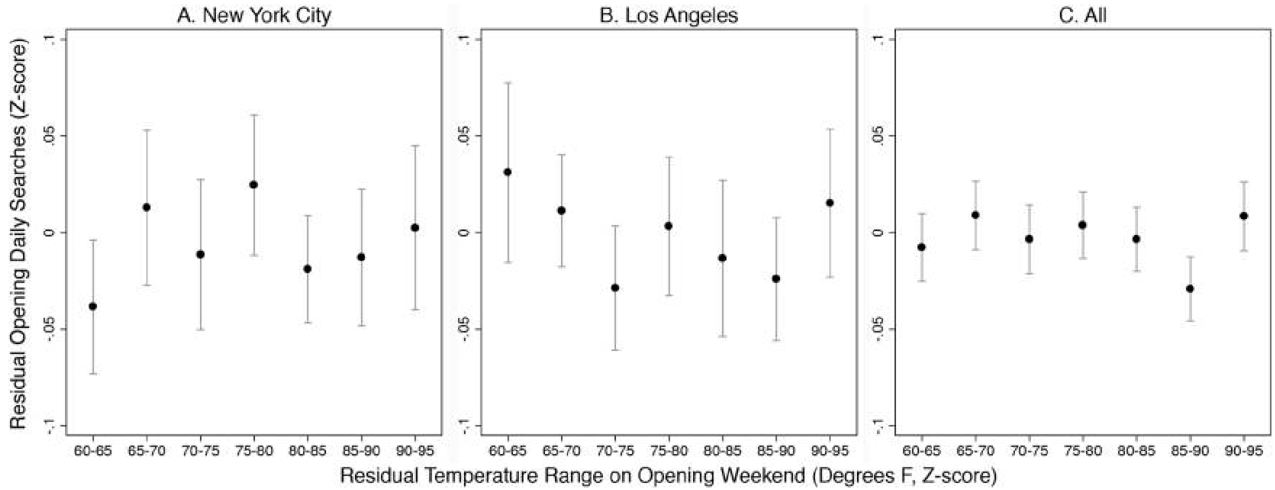

FIG. 6.—The effect of local weather shocks on local viewership. The local analogue of figure 2, each panel plots the coefficient of the regression of abnormal opening weekend local Google searches on each listed weather shock, along with the corresponding 95 percent confidence intervals, controlling for abnormal local searches in the weekend prior to opening and for movie fixed effects. Weather shocks are measured as the *Z*-score of the residual of the indicator for the MSA in each temperature range, and abnormal local searches are as described in the text. Panel A corresponds to just New York, panel B to just Los Angeles, and panel C to all five MSAs in our sample. Observations are at the movie by date by MSA level (576 observations in panel A, 480 in panel B, and 2,064 in panel C).

With our MSA-level data, the LASSO-chosen instrument is a residualized indicator for the 85°–90° temperature range.[34] Figure 7 shows the corresponding first-stage relationship in a binned scatter plot: when unexpectedly in this “good” temperature range locally, abnormal viewership locally is lower. The binned scatter plots in figure 8 show the corresponding reduced-form relationship between opening weekend weather and ticket sales in subsequent weekends: when local weather is unexpectedly good on the opening weekend, abnormal local viewership is lower in each subsequent weekend.

Table 4 reports the magnitudes of network externalities observed locally from instrumented moviegoing shocks at the local level. Here, the dependent variable is local movie searches in week w and the independent variable is local movie searches in the opening weekend, instrumented with local weather.[35] The first row includes no additional controls; the second row includes controls for movie fixed effects; and the third row, which corresponds to the figures, includes controls for movie fixed effects and for local searches in the weekend just prior to opening. The addition of movie fixed effects and prior local searches reduces the point estimates slightly (insignificantly), but in all cases the results are highly comparable to our overall estimates (table 2), suggesting that the observed momentum is driven by social spillovers arising predominantly at the local level.[36]

VI. A Role for Learning?

We have demonstrated a strong relationship between abnormal viewership of a movie in the opening weekend and abnormal viewership of that movie in subsequent weekends, even when the former was driven by exogenous shocks orthogonal to movie quality. We have also shown that the estimated momentum is not driven by supply shifts and that it is largely local in nature. The particular nature of the local demand shifts, however, remains an open question. In this section, we look for evidence that a straightforward social or observational learning story is driving our results.

[34] This is different from our main LASSO-chosen national weather instrument, likely because the impact of the weather is different across different locations; we would not expect the populations, weather preferences, and relative alternatives to moviegoing in the five largest metropolitan areas to be nationally representative.

[35] See online app. table A.5 for the corresponding first-stage results.

[36] If consumers are forward looking and weather forecasts are sufficiently accurate, searches in the weekend prior to opening might actually constitute a direct measure of opening weekend demand that takes weather shocks into account. However, controlling instead for searches 2 weeks prior does not change the results; for brevity those results are excluded here but are available on request.

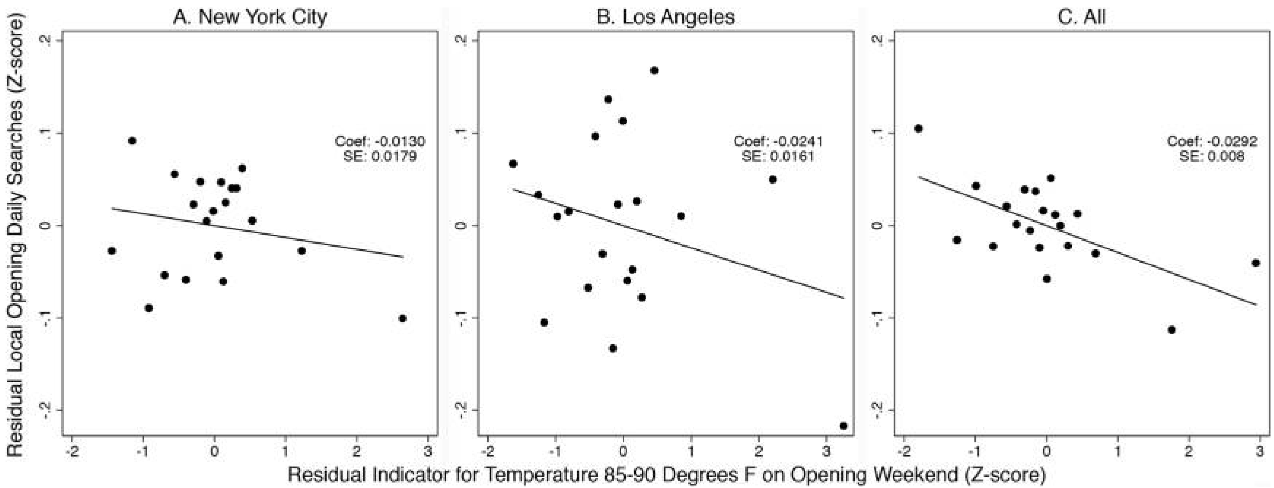

FIG. 7.—Local first-stage bin scatters. The local analogue of figure 4, each panel plots the abnormal opening weekend local searches against the residual of the indicator for whether the MSA is in the 85°–90° range, controlling for abnormal local searches in the weekend prior to opening and for movie fixed effects. For exposition, the weather shock measure is grouped into 100 equal-sized bins; each point corresponds to the mean local weather shock and local abnormal viewership within a bin. Panel A corresponds to just New York, panel B to just Los Angeles, and panel C to all five MSAs in our sample. The slope of the line of best fit and the corresponding robust standard error clustered by date are included in each panel.

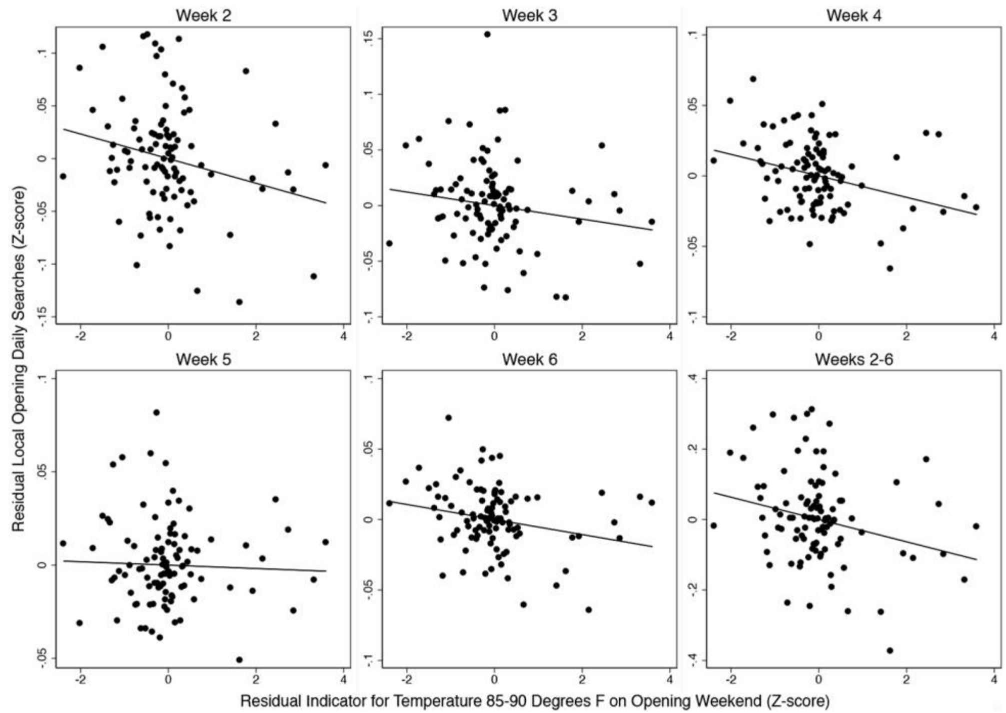

Fig. 8.—Local reduced-form bin scatters. The local analogue of figure 5, the panels plot the residual of abnormal local searches in each subsequent weekend 2–6, and overall across those weekends, against a residualized indicator for the MSA's weather being in the 85°–90° range on opening weekend, controlling for abnormal local searches in the weekend prior to opening and for movie fixed effects. For exposition, the weather shock measure is grouped into 100 equal-sized bins; each point corresponds to the mean weather shock and abnormal viewership within a bin.

TABLE 4
Local Momentum from Network Externalities

	Week 2	Week 3	Week 4	Week 5	Week 6	Weeks 2–6
Local searches	.518***	.320***	.227***	.0969***	.192***	1.354***
	(.0801)	(.0639)	(.0459)	(.0296)	(.0348)	(.215)
Controls:						
Movie fixed effects	No	No	No	No	No	No
Weekend preopening	No	No	No	No	No	No
Local searches	.347***	.181**	.226***	.0269	.159**	.939***
	(.108)	(.0922)	(.0750)	(.0512)	(.0644)	(.245)
Controls:						
Movie fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Weekend preopening	No	No	No	No	No	No
Local searches	.297**	.156	.229***	.00369	.157**	.843***
	(.126)	(.109)	(.0877)	(.0613)	(.0744)	(.283)
Controls:						
Movie fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Weekend preopening	Yes	Yes	Yes	Yes	Yes	Yes

Note.—This table shows how local searches in subsequent weekends are affected by local searches on opening weekend. The table replicates the IV results from table 2 using local Google searches as a proxy for ticket sales. The results in the first row include no additional controls; the results in the second row include controls for movie fixed effects; the results in the third row include controls for movie fixed effects and for local searches in the weekend prior to opening. The first-stage results are included in online app. table A.5. Observations are at the movie by date by MSA level (2,064 observations). Standard errors, clustered at the date level, are in parentheses. The corresponding OLS estimates are presented in online app. table A.6.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

At the most basic level, a model in which momentum is generated by some type of learning differs from a model in which momentum is generated by network externalities via the role of private information. Momentum originating from network externalities need not require consumers to have private information. In our model, for example, all consumers hold the same (unchanging) views about quality, and there is nothing for consumers to learn. All information is public, and momentum is generated by the combination of heterogeneity in preferences and network externalities: some consumers find it worthwhile to view a movie irrespective of its viewership, and these consumers view the movie as soon as it opens; meanwhile, others view the movie only when prior viewership is sufficiently high that they find it worthwhile to view the movie themselves.

Much of the existing theoretical research on crowd following, in contrast, focuses on the role of information and, in particular, on models of social learning and observational learning (see, e.g., Banerjee 1992; Bikhchandani et al. 1992, 1998; Ellison and Fudenberg 1995; McFadden and Train 1996; Çelen and Kariv 2004; Young 2009). The precise mechanisms and contexts vary, but in brief, the individual is generally assumed to have imperfect information about the quality of a good or experience and so takes into account the observed choices and/or reports of others in making her own decision. Because the existence of private information that can somehow be passed from one consumer to another creates an opportunity for learning, our strategy for distinguishing momentum due to network externalities from momentum due to learning is to focus on the role of information. Before proceeding, we note that the learning literature distinguishes between models in which information about payoffs is communicated directly—typically called "social learning"—from models in which consumers simply observe whether or not others are consuming the good and thus interpret such consumption as a signal of quality—typically called "observational learning." We focus first on a test of social learning before turning to a test of observational learning.

A. *Social Learning*

By instrumenting with shocks that are orthogonal to movie quality, we have isolated shocks to opening weekend viewership that are plausibly independent of quality. These viewership shocks should, in and of themselves, thus provide no quality signal and need not induce quality updating among individuals considering attending the movie in later weeks. Nevertheless, we might wonder whether larger early viewership boosts later sales in part because it influences the availability of information about the movie.

Disentangling a network externalities story from a social learning story is not easy. For one, a sufficiently complex model of moviegoing in the presence of network externalities could predict behavior patterns that are similar to those predicted by a learning story. Moreover, as Young (2009) notes, the theoretical literature on product adoption through learning is characterized by substantial diversity in behavioral and informational assumptions; taken together, learning models can generate a large and varied set of empirical predictions, making it hard to completely rule out a learning story.

To make progress, we focus first on an intuitive prediction of behavior in the presence of social learning (and address a separate prediction of behavior in the presence of observational learning in Sec. VI.B): in the presence of social learning, shocks to viewership should induce stronger momentum for high-quality movies than for low-quality movies. This learning prediction arises out of a canonical normal-normal model of social learning, such as that elegantly exposited by Young (2009), in which agents are risk neutral and hold private (conjugate) priors about the normally distributed value of viewing. In this model, the net utility of viewership has both a private and a public component, and social learning occurs over time as agents view movies and share their payoffs with their peers. Since weather-induced shocks to viewership on opening weekend are orthogonal to movie characteristics, such shocks simply change the amount of information available after opening weekend.

In online appendix E.1, we show formally how the effect on subsequent viewership of an exogenous shock to prior viewership varies according to movie quality. For intuition, suppose that there are two movies that are identical aside from their expected (ex ante) and realized (ex post) mean utilities of viewing; that is, viewers correctly expect that one movie is of higher quality than the other. On opening weekend, potential viewers have only their noisy priors with which to make attendance decisions; in subsequent weekends, new potential viewers form expectations by updating their priors using the additional information from those who have already viewed. As Young (2009) shows, potential viewers can be characterized by their *information thresholds*, or the amount of prior viewership required in expectation to convince a potential viewer to attend. Higher-quality movies have lower information thresholds (would-be viewers need less convincing), so all else equal, a shock to prior viewership is more likely to push a potential viewer who has not yet found it optimal to attend either movie over her information threshold for the high-quality movie than for the low-quality movie. In other words, the effect of the shock is greater for the high-quality movie than for the low-quality movie. This is a useful distinction because, as we demonstrated in Section II, momentum generated from network externalities need not depend on movie quality.[37]

As a test for social learning, then, we ask whether the momentum generated from an exogenous shock to opening weekend viewership is stronger for higher-quality movies than for lower-quality movies. We proxy for realized movie quality with ratings by expert reviewers as described in Section III. Panel A of table 5 shows estimated momentum separately by high-quality (in the top tercile) and low-quality (in the bottom tercile) movies. The results do not present compelling evidence that our estimated momentum varies with quality: relative to movies with low expert ratings, movies with high expert ratings experience about the same levels of momentum in early weeks and only slightly (insignificantly) more momentum in later weeks. However, we note that the confidence interval on the difference admits differences in momentum as large as 0.69, which is economically significant.[38]

[37] That quality influences adoption paths in a learning framework has been used in the literature in other contexts to distinguish social learning stories from other adoption models. For example, Young's (2009) analysis focuses on the features that distinguish social learning stories from social influence and contagion models, and he shows that adoption curves in learning frameworks are unusually reliant on the payoffs of adoption. Our extension of his model in online app. E.1 shows that social learning leads adoption paths of higher-payoff goods to respond more to a shock to initial adoption than lower-payoff goods.

[38] Online app. table A.7 shows the corresponding OLS results.

TABLE 5
MOMENTUM BY MOVIE QUALITY AND INFORMATION ABOUT MOVIE QUALITY

	Week 2	Week 3	Week 4	Week 5	Week 6	Weeks 2–6
	A. By Movie Quality					
High quality ($N=705$)	.484***	.292***	.208***	.140***	.119***	1.243***
	(.0842)	(.0650)	(.0501)	(.0361)	(.0308)	(.243)
Low quality ($N=825$)	.587***	.317***	.169***	.0878***	.0568***	1.217***
	(.0871)	(.0651)	(.0460)	(.0281)	(.0182)	(.214)
Difference: high − low	−.103	−.025	.039	.060*	.043*	.026
	(.120)	(.091)	(.067)	(.045)	(.034)	(.323)
	B. By Information about Movie Quality					
High information ($N=744$)	.366***	.180***	.136***	.0748***	.0543***	.812***
	(.0641)	(.0435)	(.0323)	(.0226)	(.0179)	(.167)
Low information ($N=705$)	.368***	.285***	.140***	.0563**	.0323**	.881***
	(.0606)	(.0460)	(.0283)	(.0251)	(.0164)	(.151)
Difference: high − low	−.002	−.105*	−.004	.018	.022	−.069
	(.087)	(.062)	(.042)	(.033)	(.023)	(.225)

NOTE.—Panel A replicates the IV results from table 2 separately for high- and low-quality movies, defined as the top third and bottom third in ratings by top 1,000 voters, respectively. Top 1,000 voters are the 1,000 people who have voted for the most titles in IMDb ratings polls; high rated corresponds to 6.3 and above; low rated is 5.6 and below. The final column of panel A reports the differences in the point estimates. Panel B does the same separately for movies in the top and bottom thirds by production budget, respectively, and the final column again reports the differences in the point estimates. The first-stage results are included in online app. table A.3. The corresponding OLS estimates are presented in online app. table A.7.
* Significant at the 10 percent level.
** Significant at the 5 percent level.
*** Significant at the 1 percent level.

B. *Observational Learning*

In a standard observational learning model (e.g., that set forth in Banerjee [1992] or in Bikhchandani et al. [1992]), an individual holding private beliefs about the utility of consuming a good updates her priors by looking at the previous choices made by others. In our context, observational learning predicts that the viewership decisions of others inform new potential viewers about where their priors are situated relative to those of others.

Observational learning could take place in a variety of ways. Perhaps the most plausible observational learning mechanism in moviegoing is one in which potential viewers interpret past box office success as a signal of quality. However, the results presented in Section V suggest that this mechanism is unlikely to be driving our results given that our estimated momentum is largely local in nature and box office sales are predominantly reported at the national level. Alternatively, observational learning could occur locally, for example, if individuals see the attendance decisions of their local peers (e.g., in person or through social media).

In observational learning, the precision of priors determines the extent to which prior viewership informs new potential viewers' beliefs about quality.[39] In online appendix E.2, we incorporate an exogenous shock to early adoption in the model of observational learning set forth in Bikhchandani et al. (1992) and show that a shock to early adoption generates less momentum the more precise priors about quality are (i.e., the less uncertainty there is about quality). For intuition, consider the adoption paths of two movies that have identical expected and realized qualities, but consumers have private beliefs that are precise for one movie and diffuse for the other. With observational learning, a shock to opening weekend viewership is more meaningful for the movie for which priors are diffuse because for this movie each incremental ticket sale provides a relatively stronger signal of quality. In the limit, ticket sales for a movie for which priors are perfectly precise would be unaffected by an opening weekend viewership shock.

Panel B of table 5 analyzes the role of uncertainty in our estimated momentum. As described in Section III.C, we proxy for uncertainty with production budget and report momentum estimates separately for movies that fall in the top tercile and the bottom tercile. The final row reports the differences between the point estimates. The results do not present compelling evidence that movies are differentially affected according to the level of ex ante uncertainty about quality: although low-budget movies exhibit slightly higher momentum in week 2, they actually have slightly lower momentum in all other weeks, and in no week is the difference in estimated momentum between high- and low-budget movies statistically significant.[40]

In sum, although we do not find compelling evidence that a simple model of social or observational learning is driving our estimated momentum, we cannot rule out some role for learning. Other learning models are certainly possible, and learning stories may of course coexist with network externalities (e.g., Choi 1997). Nonetheless, the close alignment between the estimates presented in this section split across both quality and the precision of information about quality suggests that the observed momentum is driven in part by a preference for shared experience, and not only by learning. In related work, Moretti (2011) analyzes learning in moviegoing and attempts to rule out the possibility that there is momentum from network externalities by instrumenting for opening weekend viewership with weather. In online appendix F, however, we show that Moretti's rejection of network externalities is not robust to the inclusion of seasonal controls in his specification.

[39] See, e.g., Bikhchandani et al. (1992) for a general formalization.
[40] Online app. table A.7 shows the corresponding OLS results.

VII. Economic Implications

We have demonstrated that the observed momentum is a demand-side phenomenon, that it is local in origin, and that it is driven at least in part by a preference for shared experience. In this section, we explore what decisions these social spillovers drive and for what types of audiences the momentum is strongest. We then discuss the broader economic implications of the observed momentum.

A. *Substitution*

From where are these marginal viewers in subsequent weeks coming? In table 6, we analyze to what extent the marginal viewers are substituting across movies versus across activities. In the first row, the endogenous regressor is abnormal viewership of new releases in week w and the outcomes are abnormal viewership in $w + 1$ of (1) those same movies (i.e., our base case results), (2) all movies showing in both w and $w + 1$, (3) new movies opening in $w + 1$, and (4) all movies showing in $w + 1$, respectively.[41] Each reported coefficient is from a separate regression. The first column simply reproduces our base case results. The second shows that (unsurprisingly) shocks to opening weekend viewership are correlated with higher viewership in $w + 1$ of all movies that played both weekends.

The first and third columns taken together provide suggestive evidence of some substitution across movies: for 100 more viewers of movies opening last weekend, we see 47 more viewers of those movies this weekend and 35 fewer viewers of new movies just opening. This is consistent with social spillovers increasing the utility of seeing movies that did particularly well last weekend, which in turn leads to an increase in overall demand for these movies and a corresponding reduction in demand for this weekend's new movies (which experienced no such shock). The positive yet statistically imprecise point estimate in the final column leaves us unable to identify conclusively whether this produces momentum in aggregate moviegoing or whether the effect is entirely inframarginal.

Since weather shocks may well engender momentum for any movie showing both this weekend and the next (not just movies that opened this weekend), the second row of table 6 shows the corresponding results when the endogenous regressor is defined as ticket sales in week w for all movies that showed in both w and $w + 1$.[42] With this specification, we again find strong momentum and some evidence of substitution away from new movies released the following weekend. For 100 additional

[41] The first stage is the same as in our base case (see table 1).

[42] The instrument is the same as in our base case specifications; the first stage is reported in online app. table A.3.

TABLE 6
Substitution across Movies and Activities

Endogenous Regressor: Audiences This Week	Outcome Variable: Audiences Next Week			
	Movies in Second Week	Movies in Second to Sixth Weeks	Movies in First Week	Movies in First to Sixth Weeks
Movies in first week	.474***	.417***	−.352*	.0482
	(.0474)	(.0820)	(.197)	(.182)
Movies in first to fifth weeks		.460***	−.388*	.0531
		(.0680)	(.213)	(.201)

Note.—Each reported coefficient is from a separate regression. In the first row, the endogenous regressor is abnormal daily ticket sales weekend w for movies that opened in weekend w; the outcome variables are abnormal daily ticket sales in weekend $w + 1$ for movies that (1) opened in week w, (2) played in both w and $w + 1$, (3) opened in week $w + 1$, and (4) played in week $w + 1$, respectively. The corresponding first stage is in the first row of table 1. In the second row, the endogenous regressor is abnormal daily ticket sales in weekend w for movies that played in both w and $w + 1$; the corresponding first-stage results are reported in online app. table A.3.
* Significant at the 10 percent level.
*** Significant at the 1 percent level.

viewers in weekend w to movies showing in both w and $w + 1$, we observe about 46 more viewers of those same movies in $w + 1$; an estimated 80 percent of these would otherwise have seen one of the new releases in $w + 1$. In sum, a positive, weather-induced shock this week leads next week's viewers to substitute away from next week's new releases and toward movies that are randomly popular in theaters this week.

B. *Social Multipliers by Age*

Recall that the model in Section II predicts that the more viewers value the viewership of others—that is, the stronger network externalities are—the larger viewership is in all subsequent weekends relative to viewership in the opening weekend. In this section, we explore the empirical implications of this prediction by splitting movies according to the age of their target audience; to the extent that different age groups place different values on peer viewership, our model predicts that we should see divergence in estimated momentum.[43]

To examine momentum by audience age, we classify each movie into one of two categories based on its age appropriateness according to the MPAA: (1) "child-friendly" and (2) "adult-oriented." Child-friendly includes all movies with a G (general audiences; all ages admitted) or PG (parental

[43] We might observe divergence in estimated momentum by age for other reasons, including but not limited to an observational learning story in which adults are the decision makers and have more diffuse priors about the quality of child-oriented movies. While we did not find compelling evidence of divergence in estimated momentum by the strength of the priors as explored in Sec. VI.B, we note here that we cannot rule out a role for learning.

guidance suggested; some material may not be suitable for children) MPAA rating; adult-oriented movies are those rated PG-13 (parents strongly cautioned; some material parents might consider inappropriate for children under 13 years) or R (restricted; people under 17 years may be admitted only if accompanied by a parent or guardian).

Table 7 shows estimated momentum separately by age suitability, and online appendix figure B.4 plots these estimated network externality effects by week in theater. Child-friendly movies exhibit significantly stronger momentum from network externalities in early weeks: for 100 additional viewers opening weekend, child-friendly movies bring in just over 70 additional viewers the second weekend, compared with just 46 additional viewers for adult-oriented movies. Although the momentum among children falls in later weeks (and by the fourth week marginal momentum in child-friendly movies is about on par with adult-oriented movies), cumulative momentum for child-friendly movies is roughly 50 percent higher than that for adult-oriented movies.

C. Aggregate Magnitudes

Altogether, our work presents two basic sets of facts: the first is that the weather has a significant influence on contemporaneous ticket sales; the second is that initial viewership shocks, such as those due to the weather, have multiplicative effects on subsequent viewership through social spillovers. In what follows, we discuss the economic implications of these effects.

1. Moviegoing and Opening Weekend Weather

We estimate that weather shocks explain an important fraction of the variation in both opening weekend ticket sales and total ticket sales. Our main measure capturing the percentage of theaters unexpectedly in the 75°–80° range opening weekend explains 1.8 percent of opening weekend ticket sales, and all of the opening weekend weather measures we provide to LASSO together explain 8.5 percent of opening weekend sales.[44] Through social spillovers, the effects of these opening weekend weather shocks continue to be realized in subsequent weeks. Defining aggregate ticket sales as all sales during the first six weekends in theaters, we find that the percentage of theaters unexpectedly in the 75°–80° range

[44] The corresponding numbers in terms of abnormal (i.e., residualized) opening weekend ticket sales are quite similar: the percentage of theaters unexpectedly in the 75°–80° range opening weekend explains 1.8 percent of opening weekend abnormal ticket sales, and all of the opening weekend weather measures we provide to LASSO together explain 9.6 percent.

TABLE 7
Network Externalities by Age Suitability

	Week 2	Week 3	Week 4	Week 5	Week 6	Weeks 2–6
Child-friendly movies						
($N = 802$)	.691***	.402***	.150***	.0610*	.0490**	1.354***
	(.116)	(.0727)	(.0468)	(.0353)	(.0243)	(.213)
Adults-only movies						
($N = 1{,}629$)	.440***	.201***	.122***	.0989***	.0791***	.941***
	(.0493)	(.0361)	(.0269)	(.0202)	(.0158)	(.132)
Differences: child − adult	.251**	.201***	.028	−.037	−.03	.413*
	(.126)	(.081)	(.054)	(.041)	(.029)	(.251)

Note.—This table replicates the IV results from table 2 separately by movie MPAA rating category. Child-friendly movies are those rated G or PG and adult-oriented movies are rated PG-13 or R. The final column reports the differences in the point estimates. The first-stage results are included in online app. table A.3.
* Significant at the 10 percent level.
** Significant at the 5 percent level.
*** Significant at the 1 percent level.

opening weekend explains 1.7 percent of aggregate sales, and all of the opening weekend weather measures we provide to LASSO together explain 10.9 percent.[45] This is sizable relative to other demand shifters. For example, Moretti (2011) finds that TV advertising explains 48 percent of aggregate ticket sales, suggesting that the effect of opening weekend weather shocks on aggregate sales (and thus aggregate revenue) is more than 20 percent that of advertising.

2. Network Externalities and Firm Incentives

The multiplier effect we estimate is consistent with the movie industry's strong focus on opening weekend viewership. This focus can be seen in the intense competition for release dates timed on major holidays and in the estimated 90 percent of advertising outlays that occur in advance of opening (Elberse and Anand 2007). Prior research has proffered explanations for the industry's focus on opening weekend outcomes that do not take into account the social multiplier we estimate. For example, Caves (2001) argues that opening weekend viewership is cheaper to acquire in expectation because it does not rely as much on movie quality. Our findings build on this literature by demonstrating that network externalities double the marginal benefit of an opening weekend viewer: if there is some reason to believe that viewership on opening weekend is cost-effective to

[45] The corresponding numbers in terms of abnormal ticket sales are again quite similar: the percentage of theaters unexpectedly in the 75°–80° range opening weekend explains 1.8 percent of aggregate abnormal ticket sales, and all of the opening weekend weather measures we provide to LASSO explain 11.1 percent.

acquire relative to viewership in other weekends, then taking the multiplier into account makes it twice as cost-effective.[46]

VIII. Conclusion

In this paper, we exploit the randomness of weather, and the relationship between weather and moviegoing, to test for and quantify network externalities in movie consumption. In the first stage, we instrument for opening weekend viewership with unanticipated and plausibly exogenous weather shocks that weekend. We implement LASSO variable selection methods to select from among the large number of potential weather instruments in order to generate a strong and econometrically sound first stage. We expect that this approach will prove similarly useful in other settings in which weather is a powerful and exogenous determinant of behavior, but specifying the appropriate first stage is otherwise nonobvious.

Using our LASSO-chosen instruments, we estimate the effect of exogenous shocks to opening weekend viewership on viewership in later weekends. Consistent with our simple model of moviegoing in the presence of network externalities, our results show that a shock to opening weekend viewership is doubled over the following five weekends. Almost all of this effect is observed at the local (MSA) level. Although we cannot reject the hypothesis that our results are at least in part generated by a learning story, the fact that our estimated momentum is largely local and that it varies with proxies for neither quality nor the precision of prior information about quality suggests that the role of learning is limited.

While this paper has focused on in-theater movies, our findings may extend to other settings in which herd behavior is observed. DellaVigna et al. (2014) note that research examining why people vote has focused on two potential explanations: (1) that individuals vote because they seek to affect the outcome of the election or (2) that individuals vote because they believe it is the right thing to do. Our work may suggest a third: that individuals vote because they know others are voting, and they value sharing in that experience. If so, policies designed to accentuate or enhance the shared experience of voting could be particularly effective in increasing voter turnout.

Network externalities could also be an implement in the toolbox of the "economist as engineer" (Roth 2003) to enhance otherwise largely independent experiences. With the rise of relatively solitary activities such as gaming, remote work, and online learning, we expect that further research into where and how platforms might leverage the positive effects

[46] Although our estimates do not speak to whether the social multiplier is largest when it arises out of viewership on the opening weekend vs. out of viewership on subsequent weekends, our work shows that there is a large (opening weekend) social multiplier, so that the aggregate contribution of any given demand shifter is larger than its contemporaneous effect.

of network externalities to deepen participation and engagement would prove fruitful.

Appendix A

Holiday Controls

Our holiday indicators are exactly those of Dahl and DellaVigna (2009) and are similarly motivated by the fact that (1) holidays have an impact on movie audience sizes (usually positively), (2) the effect varies across holidays, and (3) audience sizes are often also affected in the days just around each holiday. We include indicators for Martin Luther King Day, Presidents' Day, Memorial Day, Labor Day, Columbus Day, Independence Day, Veterans Day, Easter, Thanksgiving Day, Christmas Eve, Christmas, New Year's Eve, New Year's Day, Saint Patrick's Day, Valentine's Day, Halloween, Cinco de Mayo, and Mother's Day. We also include separate indicators for the Friday, Saturday, and Sunday before each of MLK Day, Presidents' Day, Memorial Day, Labor Day, and Columbus Day; for the Friday and Saturday before Easter; for the Wednesday before Thanksgiving and for the weekend after; for the four days before Christmas Eve (December 20–23) and the five days after Christmas (December 26–30); and for the two days after New Year's (January 2–3). Finally, for Independence Day, Veterans Day, Christmas, New Year's, and Valentine's Day, we include an indicator for whether each falls on a Saturday or Sunday. Several of these indicators drop out when we restrict our sample to movie weekends (Friday, Saturday, Sunday) only.

Appendix B

Google Trends Search Data as a Proxy for Viewership

We use the Google Trends search data at the day by MSA by topic level, the most granular level at which it is made publicly available. Although Google Trends data are available for specific queries, we use the topic classification engine, which classifies searches as pertaining to particular movies.[47]

The raw data consist of integral figures ranging from 0 to 100, where 0 and 100 are the lowest and highest points, respectively, in any single data export. Daily data can be exported only in 3-month windows; and if volumes are sufficiently low for the entire period of intended export, then daily data are not available. In this case, Google provides weekly data or, if volumes still remain too low, monthly data; but even weekly data are insufficient for our analysis because the weeks are measured from Sunday to Saturday, which does not allow us to distinguish between weekends since release.

[47] Without a topic classification system, it can be difficult to determine which queries relate to specific movies. For example, a simple search for "Superman" could pertain to one of the many Superman movies, comic books, or other Superman merchandise. Thus, we use Google's topic classification service to classify searches as pertaining to specific movies, such as the 2006 film *Superman Returns*.

Since Google censors to 0 any observation for which the total search volume falls below some (undisclosed) threshold, we are restricted in the number of movies and cities for which we can undertake the local analysis of network externalities.[48] We begin with the top 500 movies by US gross ticket sales that were released May–September 2004–13.[49]

For each movie, we collect local search volumes in each of the 10 largest MSAs (as well as national search volumes) on each day in a 3-month period beginning 2 weeks before the movie's release. Of the 5,000 MSA by movie combinations, nearly 4,000 are censored, leaving 1,000 MSA by movie combinations that come disproportionately from the largest MSAs.[50] Our analysis is conditional on a rich set of fixed effects, so it is important that we observe search volumes for the same weekend in the same city over multiple years; given this, we restrict our data set to the five MSAs with the most observations: New York, Los Angeles, Chicago, Washington, DC, and San Francisco. This corresponds to 67 percent of the original MSA by movie sample.

Since the data are normalized so that the highest point in any given data export is 100 and data exports are limited to four topics for a 90-day period, we separately export search volume for each movie in combination with the "Harvard University" topic. We then use the trend in Harvard searches to standardize movie search data across time.[51] Finally, we convert our search measure to the Z-score of search volume within each MSA. Note that because Google provides only unitless search figures, we are unable to directly compare search volumes across MSAs.

[48] While Stephens-Davidowitz (2014) develops an algorithm to circumvent some of Google's censoring, the algorithm relies on the fact that his study is focused on search volumes relating to specific queries. Our data exports, meanwhile, rely on Google's topic classification system, which renders the algorithm ineffectual.

[49] We restrict to the top 500 movies because Google does not provide an application program interface for data access, so collection of large amounts of data from Google Trends is cumbersome. We focus on May through September because the summer release season is associated with significant variance in the instrument of choice in our main analysis (the percentage of establishments with maximum temperature unexpectedly between 75° and 80°); we do not include 2014 because at the time of writing Google's topic classification of searches pertaining to movies released in 2014 was largely incomplete.

[50] The extent of the censoring and the selection it might induce appear to be largely related to idiosyncratic factors affecting the extent to which Google's topic classification system is able to successfully classify searches. For example, *Ironman*, which grossed over $300 million in the United States, is censored throughout, while *Ironman 2*, which grossed nearly as much, is not.

[51] We use Harvard-related searches as an aggregator because searches for Harvard always show substantial volume, while individual movies tend to have nonzero volume only in the weeks surrounding release.

References

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *Q.J.E.* 107 (3): 797–817.

Becker, Gary S. 1991. "A Note on Restaurant Pricing and Other Examples of Social Influences on Price." *J.P.E.* 99 (5): 1109–16.

Belloni, Alexandre, Daniel Chen, Victor Chernozhukov, and Christian Hansen. 2012. "Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain." *Econometrica* 80 (6): 2369–2429.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2011. "LASSO Methods for Gaussian Instrumental Variables Models." https://arxiv.org/abs/1012.1297.

———. 2014. "Inference on Treatment Effects after Selection amongst High-Dimensional Controls." *Rev. Econ. Studies* 81 (2): 608–50.

Belloni, Alexandre, Victor Chernozhukov, Christian Hansen, and Damian Kozbur. 2014. "Inference in High Dimensional Panel Models with an Application to Gun Control." https://arxiv.org/abs/1411.6507.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *J.P.E.* 100 (5): 992–1026.

———. 1998. "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades." *J. Econ. Perspectives* 12 (3): 151–70.

Bursztyn, Leonardo, Florian Ederer, Bruno Ferman, and Noam Yuchtman. 2014. "Understanding Mechanisms Underlying Peer Effects: Evidence from a Field Experiment on Financial Decisions." *Econometrica* 82 (4): 1273–1301.

Cabral, Luis, and Gabriel Natividad. 2016. "Box-Office Demand: The Importance of Being #1." *J. Indus. Econ.* 64 (2): 277–94.

Cattaneo, Matias D., Michael Jansson, and Whitney K. Newey. 2015. "Alternative Asymptotics and the Partially Linear Model with Many Regressors." Manuscript, Univ. California, Berkeley.

Caves, Richard E. 2001. *Creative Industries: Contracts between Art and Commerce*. Cambridge, MA: Harvard Univ. Press.

Çelen, Boaçhan, and Shachar Kariv. 2004. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *A.E.R.* 94 (3): 484–98.

Chen, Daniel L., and Jasmin Sethi. 2012. "Insiders and Outsiders: Does Forbidding Sexual Harassment Exacerbate Gender Inequality?" Manuscript, Toulouse Inst. Advanced Study.

Chen, Daniel L., and Susan Yeh. 2012. "Growth under the Shadow of Expropriation? The Economic Impacts of Eminent Domain." Manuscript, Toulouse Inst. Advanced Study.

Chen, Yi-Fen. 2008. "Herd Behavior in Purchasing Books Online." *Computers in Human Behavior* 24 (5): 1977–92.

Chernozhukov, Victor, and Christian Hansen. 2013. "Econometrics of High-Dimensional Sparse Models." Lecture, NBER, Cambridge, MA.

Choi, Jay Pil. 1997. "Herd Behavior, the 'Penguin Effect,' and the Suppression of Informational Diffusion: An Analysis of Informational Externalities and Payoff Interdependency." *Rand J. Econ.* 28 (3): 407–25.

Conover, W. J. 1999. *Practical Nonparametric Statistics*. 3rd ed. New York: Wiley.

Corts, Kenneth S. 2001. "The Strategic Effects of Vertical Market Structure: Common Agency and Divisionalization in the US Motion Picture Industry." *J. Econ. and Management Strategy* 10 (4): 509–28.

Dahl, Gordon, and Stefano DellaVigna. 2009. "Does Movie Violence Increase Violent Crime?" *Q.J.E.* 124 (2): 677–734.

DellaVigna, Stefano, John A. List, Ulrike Malmendier, and Gautam Rao. 2014. "Voting to Tell Others." Working Paper no. 19832, NBER, Cambridge, MA.

Donihue, Michael R., Randy A. Nelson, Donald M. Waldman, and Calbraith Wheaton. 2001. "What's an Oscar Worth?" *Econ. Inquiry* 39 (1): 1–6.

Einav, Liran. 2007. "Seasonality in the US Motion Picture Industry." *Rand J. Econ.* 38 (1): 127–45.

Elberse, Anita, and Bharat Anand. 2007. "The Effectiveness of Pre-release Advertising for Motion Pictures: An Empirical Investigation Using a Simulated Market." *Information Econ. and Policy* 19 (3): 319–43.

Eliashberg, Jehoshua, and Mohanbir S. Sawhney. 1996. "A Parsimonious Model for Forecasting Gross Box-Office Revenues of Motion Pictures." *Marketing Sci.* 15 (2): 113–31.

Ellison, Glenn, and Drew Fudenberg. 1993. "Rules of Thumb for Social Learning." *J.P.E.* 101:612–43.

———. 1995. "Word-of-Mouth Communication and Social Learning." *Q.J.E.* 110 (1): 93–125.

Farrell, Joseph, and Paul Klemperer. 2007. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." In *Handbook of Industrial Organization*, vol. 3, edited by Mark Armstrong and Robert Porter, 1967–2072. Amsterdam: North-Holland.

Frick, Mira, and Yuhta Ishii. 2016. "Innovation Adoption by Forward-Looking Social Learners." Manuscript, Yale Univ.

Jensen, Richard. 1982. "Adoption and Diffusion of an Innovation of Uncertain Profitability." *J. Econ. Theory* 27 (1): 182–93.

Katz, Michael L., and Carl Shapiro. 1986. "Technology Adoption in the Presence of Network Externalities." *J.P.E.* 94 (4): 822–41.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Rev. Econ. Studies* 60 (3): 531–42.

McFadden, Daniel L., and Kenneth E. Train. 1996. "Consumers' Evaluation of New Products: Learning from Self and Others." *J.P.E.* 104 (4): 683–703.

Moretti, Enrico. 2011. "Social Learning and Peer Effects in Consumption: Evidence from Movie Sales." *Rev. Econ. Studies* 78 (1): 356–93.

Motiere, Lionel, and James G. Mulligan. 1994. "The Market for First-Run Motion Pictures." Manuscript, Coll. Bus. and Econ., Univ. Delaware.

Moul, Charles C. 2007. "Measuring Word of Mouth's Impact on Theatrical Movie Admissions." *J. Econ. and Management Strategy* 16 (4): 859–92.

Munshi, Kaivan, and Jacques Myaux. 2006. "Social Norms and the Fertility Transition." *J. Development Econ.* 80 (1): 1–38.

Prag, Jay, and James Casavant. 1994. "An Empirical Study of the Determinants of Revenues and Marketing Expenditures in the Motion Picture Industry." *J. Cultural Econ.* 18 (3): 217–35.

Roth, Alvin E. 2003. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica* 70 (4): 1341–78.

Scharfstein, David S., and Jeremy C. Stein. 1990. "Herd Behavior and Investment." *A.E.R.* 80 (3): 465–79.

Smith, Michael D., Seth Stephens-Davidowitz, and Hal Varian. 2015. "Super Returns? The Effects of Ads on Product Demand." Manuscript, Univ. California, Berkeley.

Sorensen, Alan T. 2007. "Bestseller Lists and Product Variety." *J. Indus. Econ.* 55 (4): 715–38.

Stephens-Davidowitz, Seth. 2013a. "Unreported Victims of an Economic Downturn." Manuscript, http://sethsd.com/s/childabusepaper13.pdf.

———. 2013b. "Who Will Vote? Ask Google." Manuscript, http://sethsd.com/s/election-forecasts-using-google9-vb54.pdf.

———. 2014. "The Cost of Racial Animus on a Black Candidate: Evidence Using Google Search Data." *J. Public Econ.* 118:26–40.

Welch, Ivo. 1992. "Sequential Sales, Learning, and Cascades." *J. Finance* 47 (2): 695–732.

Young, H. Peyton. 2009. "Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning." *A.E.R.* 99 (5): 1899–1924.