Paper to be presented at
the DRUID16 20th Anniversary Conference
Copenhagen, June 13-15, 2016

Innovation Diffusion Through Generation Cohorts

Anil R Doshi
UCL School of Management
Strategy and Entrepreneurship
anil.doshi@ucl.ac.uk

Abstract

This paper focuses on the process of innovation diffusion through a population where sets of new members---called generation cohorts---enter the population over time. In certain cases where diffusion in a growing population does not follow the prototypical logistic pattern, disaggregating it by generation cohorts makes the diffusion process more apparent. Subsequent generation cohorts experience increased rates of diffusion and higher levels of early adoption. Moreover, the diffusion experienced by the existing population affects the diffusion path of later generation cohorts. After documenting the effect of generation cohorts on diffusion, strategic implications are considered, with incumbents and entrants competing to capture succeeding generation cohorts that each experiences its own path of diffusion.

Jelcodes:O33,D81

Innovation Diffusion Through Generation Cohorts

Anil Doshi*

May 29, 2016
Preliminary and incomplete. Please do not cite.

Abstract

This paper focuses on the process of innovation diffusion through a population where sets of new members—called generation cohorts—enter the population over time. In certain cases where diffusion in a growing population does not follow the prototypical logistic pattern, disaggregating it by generation cohorts makes the diffusion process more apparent. Subsequent generation cohorts experience increased rates of diffusion and higher levels of early adoption. Moreover, the diffusion experienced by the existing population affects the diffusion path of later generation cohorts. After documenting the effect of generation cohorts on diffusion, strategic implications are considered, with incumbents and entrants competing to capture succeeding generation cohorts that each experiences its own path of diffusion.

*UCL School of Management, anil.doshi@ucl.ac.uk. I would like to thank Scott Stern, Dennis Yao, Lee Fleming, and Feng Zhu for their guidance in developing this project. I also thank Sen Chai, Sharat Raghavan, and the participants of the UCL reading seminar for their feedback on this paper.

1 Introduction

The process by which an innovation diffuses through a population is important for economic growth and social progress. Because of its importance, diffusion through many kinds of populations, such as individuals in a firm (Tucker, 2008), organizations (Saloner and Shepard, 1995), local regions (Griliches, 1957), and countries (Comin and Hobijn, 2004), have been studied. In the large majority of these studies, the population is considered to be fixed. In cases when population growth is considered, it is adjusted for (on a per capita basis, for example) so as to normalize the population in any given period.

However, a group of adopters concurrently joining the population may differ from already existing members. The new group may exhibit different behaviors or enter in a different environment from prior groups. Those differences, in turn, may affect the patterns of diffusion within the newer group. Understanding how diffusion changes within a single population that is growing over time is the focus of this paper. Specifically, I investigate the diffusion of an innovation in a population where members enter that population in separate generation cohorts over time.

How might subsequent cohorts of members entering into a population be considered, relative to the existing members of that population, and what are the implications of that relationship for innovation diffusion? From one perspective, new members are, by definition, the same as existing members in terms of the characteristics that identify them as being in the same population. From this perspective, one might argue that the timing of their arrival into the population may not have any impact on diffusion at all. New cohorts may continue adopting along the same path as the existing population, effectively being subsumed by the existing diffusion process. In other words, members of new cohorts may jump to the equilibrium diffusion levels achieved by the existing population.

Alternatively, new cohorts do have some differences from the existing population. Members of a new cohort join the population as a distinct group from the existing population—their adoption decisions are thus within the context of the timing of their arrival into the population (Nelson and Winter, 1982). Moreover, a newly entering cohort of members has the benefit of observing the experiences of the existing population with respect to the behaviors of adoption and extent of diffusion. Acting as a separate group with the vicarious experience of prior cohorts might indicate that each cohort would experience its own path of diffusion, resulting in successive diffusion curves for each cohort.

In this paper, I reconcile these different perspectives and argue that with respect to diffusion, generation cohorts act partly like the existing population and partly like their own distinct group of adopters. Cohorts arrive at different times and thus follow their own path of diffusion, as if they were separate populations. But as later cohorts are able to observe the pattern of diffusion into the existing population, those later paths may more quickly converge on the population equilibrium diffusion path. The result is a change in the shape of diffusion by subsequent cohorts.

To develop this argument, I theorize the impact of generational entry into a population on the diffusion of a technology in to that population. I then document its occurrence into the context of the diffusion of a social media technology into the television industry. Specifically, I look at how subsequent premiere years of television shows adopt Twitter, a social media platform. I model the diffusion for each generation cohort of newly premiering television shows and show two main results. First, later cohorts experience a shift of the diffusion curve. Second, the rate of adoption increases for subsequent cohorts in more recent periods.

Finally, I test whether prior adoption impacts subsequent generation cohort diffusion. I show that diffusion behavior through the existing population is a positive determinant of subsequent cohort diffusion, using a variety of measures. The results, when taken together, imply that diffusion across the existing population resolves some uncertainty in later cohorts, allowing for faster convergence to equilibrium diffusion rates. In the case of a more widely diffused technology, that would mean higher levels of initial adoption and faster rates of adoption.

Identifying and describing the impact of innovation diffusion through generation cohorts has important theoretical and practical implications. First, generation cohorts occur in a wide number of settings and exist in whenever the target population of an innovation is not fixed over time. Consider incoming classes of students to a high school or university. Technologies and practices implemented by the school must diffuse through each generation of first year students. Moreover, any technology that imposes age restrictions will experience generation cohorts entering in subsequent periods over time, as new prospective adopters meet the requirements. Many other online services have age minimums that would create an influx of new prospective users each year. With respect to firms, employee hiring implies that technologies being used by a firm (Tucker, 2008) might diffuse differently through waves of new employees. Within an organization, new features or safety technologies diffuse through new makes of automobiles produced each year. At an industry level, if firms are adopting a technology, then firm entry (Gort and Klepper, 1982) represents new members of a population. Individual migration and organizational entry (Alcacer and Chung, 2007) into a new geography over time may also represent generation cohorts within that region if the population is geographically bound. In general, anytime there is entry into a population, there will exist generation cohorts that will require accounting for when understanding the diffusion process.

Second, generation cohorts may be a factor in resolving why some technologies do not diffuse in a prototypical s-curve pattern through certain populations (Comin et al., 2006). In some of those cases, generation cohorts into the population may be causing the deviations. After accounting for generation cohorts in a population, the underlying diffusion process would more easily observed in those cases.

Third, if a population experiences a generation cohort effect, there are implications for the traditional model of population segmentation among adopters (Rogers, 1995). Since the timing of arrival of generation cohorts is typically exogenous to the adoption decision of a given technology, there is likely to be a similar distribution of the traditional categories of adopters. If diffusion in later cohorts is affected diffusion within the existing population, then traditional categories may be extended or compressed in terms of the timing of their adoption.

Fourth, there are strategic implications if diffusion is distinct for each generation cohort. Understanding how the process of diffusion may be different from, yet dependent on, the existing population will help firms facilitate the penetration of their innovations into the adopting population. Separately, entrants may differentiate a technology for a new generation cohort in hopes of replacing an incumbent. Incumbents would, in turn, need to consider how to deploy their technology for rapid diffusion in new generations.

This paper makes three primary contributions to the literature on the diffusion of innovations. First, I identify generations of a population as an important source of heterogeneity in a population that affects the diffusion process. Groups, or generations, of the population can enter into the population over time, which results in a series of separate, but related, diffusion processes. Diffusion patterns of successive generations differ from and are affected by prior generations. Whereas much of the diffusion literature assumes a fixed population that an innovation diffuses through or investigates cross-sectional heterogeneity in a population, I focus on how the population can change over time and the impact that has on diffusion. Second, I show that even for an innovation where actual costs of adoption are low and the innovation is a standard for the population in question, there is still a process of diffusion that takes place for subsequent cohorts of the population.

2 Diffusion and Adoption

Innovation diffusion theory has developed into two broad streams, consisting of 'macro' diffusion, where the rate of diffusion through a population is modeled and looks like a logistic or bounded exponential curve (Griliches, 1957; Geroski, 2000), and 'micro' adoption, where the adoption decisions of individual agents is investigated, often in the context of networks or network effects (Saloner and Shepard, 1995; Majumdar and Venkataraman, 1998; Tucker, 2008) or characteristics that influence the decision to adopt (Mansfield, 1963; Hannan and Freeman, 1984; Caselli and Coleman, 2001).

Across both streams, the impact of varying kinds of heterogeneity have been investigated. First is heterogeneity in the diffusing technology, specifically when there are new versions diffusing through the same population. In addition to modeling the general process of new versions of a technology (Norton and Bass, 1987), Bresnahan and Yin (2005) show that distribution of a complement technology is more important than feature improvements for the diffusion of the new versions.

Second, and more related to this paper is the heterogeneity in the population. At a national level, cross-country diffusion of a number of technologies have shown that technologies diffuse in advanced economies and then developing economies (Comin and Hobijn, 2004) and the rate of diffusion has increased for newer technologies (Comin and Hobijn, 2010). For computers, counties with highly educated populations and trade with advanced economies experienced more rapid diffusion (Caselli and Coleman, 2001).

Other research has looked primarily at the cross-sectional heterogeneity in a population. In a general model of diffusion with heterogeneous populations, Young (2009) theoretically distinguishes between contagion, social influence, and social learning.

Adopter heterogeneity has also been modeled at a micro level, with Abrahamson and Rosenkopf (1997) running simulations of diffusion through networks where nodes differ in their number of links. When adoption is considered as the intensive margin of adoption (i.e. how much an innovation is used, rather than merely a binary adoption choice), Grajek and Kretschmer (2009) find that the highest intensity users adopt earliest and later adopters are less intensive users.

One paper considers how diffusion changes over time among sets of users is Goel et al. (2015), who observe the diffusion of social media messages through subsequent waves of users in the same population and show the diversity of diffusion patterns of broadcast messages on a social media site. The population in their paper is fixed; what varies over time is the propagation of a news broadcast to succeeding levels of users through the network. The heterogeneity in the structure of the network causes the news to diffuse differentially over time.

This paper builds on the tradition of exploring diffusion at a macro level. Prior work has primarily looked at cross-sectional heterogeneity within the same population or at different populations Rather than accounting for various kinds of cross-sectional heterogeneity, in this paper, I focus on temporal heterogeneity, with the arrival of members into a population. Given the importance of understanding the process of diffusion to improved performance and economic outcomes and growth and development, it continues to be important to understand how technologies spread through populations. This paper is the first attempt to my knowledge to show how differences in a population that arise temporally impact the diffusion process of the same technology.

3 How Innovations Diffuse Through Generation Cohorts

I provide a descriptive model to show how diffusion changes across generation cohorts. The underlying intuition is based on the fact that the most recent cohort can observe the pattern of diffusion by the existing population. If the existing population has not embraced the innovation, subsequent cohorts will undergo the diffusion process as if it were a new population. Alternatively, the more the existing population has embraced the innovation, the less uncertainty that remains and the diffusion process is accelerated for the new generation cohort.

Each prospective adopter decides to adopt the innovation when the utility of adopt is greater than the outside option, where the utility is the benefits of adopting less the costs of adoption. Benefits include the value of using the adoption and, when present, its network effects (Farrell and Saloner, 1986). Costs include the actual costs of adoption and the uncertainty surrounding the innovation. Uncertainty may be technological (i.e. does this technology meet our requirements for the task and does the technology work?) or economic (i.e. will this technology become the dominant design (Anderson and Tushman, 1990), or standard (Katz and Shapiro, 1985; Farrell and Saloner, 1985), that is the norm among peer adopters (Rohlfs and Hal, 2003)?). Both perceived benefits and costs of subsequent generation cohorts, particularly uncertainty, are affected by the adoption behavior of the existing population.

In the absence of extensive diffusion within the existing population, uncertainty persists for the new cohort. Moreover, with less diffusion in the existing population, any network effects will not exert a significant impact on the adoption decision of members in the new generation cohort. Both effects imply that for new generation cohorts entering the population at a time when there is little existing diffusion, a new diffusion process takes place ensue.

However, more adoption by the existing population reduces the technological and economic uncertainty of the innovation and makes a prospective adopter in the new cohort more likely to adopt, on average. Greater adoption in the existing population also increases the propensity to adopt, by way of social influence (Bass, 1969; Easingwood et al., 1983) (observing the decision of others to adopt ensures that this innovation is the one to adopt) or social learning (observing the experience or outcomes other adopters have with the innovation) (Ellison and Fudenberg, 1993; Young, 2009). Moreover, if the innovation is one with network effects, then the importance of other users adopting is amplified by increasing the benefit to adopt. In this case, existing population diffusion effectively acts as a substitute for some part of the diffusion process of later generation cohorts, accelerating the diffusion process in that later cohort. Acceleration would be occur in either the initial adoption level or the rate of diffusion of the new generation cohort.

3.1 Initial Adoption

Diffusion through a fixed population begins among the members who derive the greatest net benefit from the technology, described in the adopter categorization framework as innovators (Rogers, 1995). Innovators are followed by early adopters, early majority, lade majority, and finally laggards.

In the presence of generation cohorts, a new generation cohort contains the same distribution of adopter categories as the existing population if the assumption that characteristics associated with a certain adopter category are orthogonal to the timing of arrival into the population. As adoption by the existing population grows, uncertainty of the innovation is reduced for all members of the new cohort. As the process of diffusion more widely takes place in the existing population, then much of the technological and economic uncertainty has been alleviated for subsequent generations. Prior cohorts have addressed the questions of whether the innovation is technically appropriate for the population. More extensive adoption within the existing population also establishes the innovation the dominant design (Anderson and Tushman, 1990).

In the new generation cohort, a larger proportion of innovators—who would have adopted over time in the initial period had they existed in the population when the technology was introduced—all experience positive net benefits from the innovation based on the reduced uncertainty provided by the existing population. Thus, innovators within the new generation cohort adopt instantly upon their arrival into the population. The more widely the technology has diffused into the existing population, the more likely later categories would experience the same rush to adopt as early innovators. Observations of the existing population effectively compress the timing of adoption of adopter categories and the compression leads to a larger level of initial adopters.

Hypothesis 1 As an innovation becomes more diffused in the existing population, initial adoption within the next generation cohort will increase.

3.2 Rate of Diffusion

As a technology diffuses through an existing population, those members present themselves as sources of diffusion through a new generation cohort, by way of social influence or social learning. The presence of a larger number of adopters of the technology for a later generation cohort suggests that the rate of diffusion would be faster for that cohort, compared with when there was a smaller existing population, as would be the case for earlier generation cohorts. As a result, later generation cohorts would experience a faster rate of diffusion as the technology becomes more widely diffused in the existing population.

Two further conditions would strengthen the increase in diffusion rate for later generation cohorts. First, if, as H1 predicts, initial levels of members of a new generation cohort, that initial mass would further speed the diffusion process among the remaining members of the cohort. In this case, increases in initial levels of adoption in later cohorts would lead to increasingly higher rates of adoption as well. Second, to the extent network effects exist in the technology, the wider diffusion in the existing population and the initial adopters in the focal generation cohort would further accelerate the rate of diffusion.

Hypothesis 2 As an innovation becomes more diffused in the existing population, the rate of adoption by the current generation cohort will increase.

3.3 Visualizing the Theory

The implications of the two hypotheses on diffusion within new generation cohorts are illustrated in Figure 1. Each panel in the figure show how diffusion might (or might not) change, depending on whether initial level and rate of diffusion for new cohorts are affected by the existing population. For the sake of simplicity, the diffusion of the existing population is represented by one initial cohort (cohort 1) and the new generation cohort is represented by cohort 2. If the extent of diffusion within existing population does not mitigate the uncertainty faced by the next generation cohort, it will follow its own path of diffusion, as in Panel A. The diffusion process would take place as if there were no prior cohorts. This would also be the result if neither hypothesis holds. The technology would diffuse through subsequent cohorts separately and independently.

If, however, diffusion into the existing population resolves some of the uncertainty around the innovation, then adoption would look like one of the panels from B through F. If only H1 holds but H2 does not, then the diffusion might appear to follow the pattern shown in Panels B or C. In Panel B, a logistic diffusion process follows a higher level of initial diffusion. Or, as Panel C shows, initial adoption might shift to a higher level, and the rate of diffusion follows the existing population. In this case, the generation cohort effectively jumps to the existing population's equilibrium diffusion path. If, however, only H2 holds, then the cohort would experience logistic diffusion with a faster rate, as shown by the steeper slope in Panel D.

If both H1 and H2, and subsequent cohorts experience higher initial levels of adoption and faster rates of diffusion, then the diffusion might look like Panel E or F. Diffusion might still be logistic, but with higher initial adoption and a steeper slope as shown in Panel E. Alternatively, after a higher level of initial adoption, diffusion may skip the beginning of the s-curve, where diffusion is ramping up, and quickly move to the equilibrium rate of diffusion within the population. In this case, diffusion for later generation cohorts might resemble a truncated logistic, or bounded exponential, curve.

4 Setting and Data

To model the diffusion through generation cohorts of a population, I use the diffusion of Twitter through newly premiering television shows each year. The diffusing technology, Twitter, is a communications platform that was launched in July 2006. The platform allows users (individuals or organizations) to publicly broadcast short messages (called Tweets) and monitor a list of selected other users' broadcasts in what resembles a newswire.

The diffusion of Twitter into the television industry was chosen because both the industry and the technology possess characteristics that make them ideally suited for this diffusion study. First, television shows represent a well-defined and complete population of at-risk members. By using the full population, no systematic bias can arise from the identification of members of the population based on the adoption decision. Television shows also have observable and systematic generation cohorts, with new television shows consistently premiering each year.

Second, as an online technology, Twitter does not release distinct versions requiring adopters to switch or upgrade (Bresnahan and Yin, 2005; Adner and Kapoor, 2015), thus ensuring that a single technology is being considered over the adoption period. Also, during the study period, there were no comparable or competing communication technologies that were being widely considered by the industry.So the diffusion decision is not affected by any considerations of switching technologies (to new versions or competing alternatives), and thus remains constant (i.e. to adopt Twitter or not) for all shows over the study period.

Third, the adoption decision by television shows of Twitter is well-defined and observable. Adoption is defined as the date when a Twitter account was created for the television show.

4.1 Data and Sample Construction

The set of television shows comes from the online database, The Futon Critic. I narrow the total set of shows to include only those shows that premiered after January 2005 and remained on air until the end of the study period in December 2013.[1] I

[1] In Section 7.2, I show that the results are largely consistent after including shows that are canceled during the study period.

then matched those shows to Twitter accounts if one was available as of the time of data collection. I requested three laborers from Amazon mTurk's service to search for the Twitter account associated with the show. If each provided a response and it was the same, I did not perform a secondary check. For each show that did not have a unanimous account identified by the mTurk laborers, and for shows that had no account returned, I performed a search for the account to reconcile the difference or to confirm the lack of an account. After collecting the account names, I used the Twitter Application Programming Interface (API) to acquire the date of account creation for all the television show Twitter profiles. The account creation date is used as the adoption date of Twitter for that television show.

A generation cohort for the sample is defined as all shows that premiered in a given year. To arrive at the sample, I aggregate show premiere and Twitter adoption (if adoption occurred) dates up to the generation cohort level, in order to observe monthly diffusion of Twitter through television shows that premiered in the same year. An observation in the resulting panel dataset is a generation cohort-month for eight generation cohorts. The first generation cohort includes shows from January 2005 to July 15, 2006—Twitter's launch date—in a single cohort (called *Pre-Twitter*). The second cohort includes shows that premiered from July 15, 2006 to December 2007 (called *July 2006-2007*). Each subsequent cohort consists of shows that premiered in each calendar year from 2008 to 2013.

Table 1 provides summary information for each of the eight generation cohorts. Diffusion is measured as the cumulative proportion of all shows in the generation cohort that have Twitter accounts as of the month of observation. Maximum diffusion (defined as the cumulative proportion of all shows that adopted twitter at the time) for the cohorts range from 0.35 to 0.47, with the exception of the Pre-Twitter cohort where adoption of shows that remained on air through the sample period reached 0.59. Second, shows adopted Twitter increasingly earlier. In fact, starting in 2010, adoptions took place prior to the year in which shows began to air.[2] Third, as a result of this 'pre-airing' adoption, initial levels of adoption were non-zero at the time shows in the generation cohort began to air. Initial levels of adoption reached as high as 0.19 in 2013.

Tables 2 and 3 provide summary statistics and correlations for the sample. There are a total of 519 months across the eight generation cohorts in the sample. The average diffusion across all months is 0.24 and 73% of months take place after shows have begun airing across all the generation cohorts. With respect to the measures on prior cohort experience, an average of 632 ($\exp(6.45)$) shows adopted in all prior years in all prior cohorts, 43 ($\exp(3.77)$) shows adopted in all prior cohorts in the prior period, and the rate of adoption among all shows in the prior period was 28% ($\exp(-1.28)$).

5 Empirical Approach

To distinguish among the diffusion curves of the different generation cohorts, I begin with the logistic function, a standard approach in the diffusion literature (Griliches,

[2]This feature of the data is given explicit consideration in the empirical specifications described in Section 5.

1957; Geroski, 2000). For a given generation cohort, i, at time, t, the general model is:

$$s_{it} = \frac{S_i}{1 + \exp\left[-(\beta_0 + \beta_1 X_t)\right]} \tag{1}$$

where diffusion, s_{it}, is the share of total shows in generation cohort i that have adopted Twitter at time t:

$$s_{it} = \frac{\texttt{accounts}_{it}}{\texttt{shows}_i}$$

Next, S_i is the maximum Twitter adoption rate for cohort i, and X_t indicates time measured in calendar months. The parameters β_0 and β_1 reflect the shift and the steepness of the curve, respectively.

To apply the above model to the context of Twitter adoption by television shows, one important adjustment has to be made, due to the fact that Twitter adoptions in later cohorts begin prior to the airing of television shows in that cohort. An empirical specification needs to account for this pre-airing 'leakage' in diffusion prior to the start of the generation cohort. To account for differences in adoption timing and rates across cohorts, as well as for differences in rates of diffusion prior to and after the start of the focal cohort, I estimate the following linear model separately for each cohort:

$$\ln\left(\frac{s_{it}}{S_i - s_{it}}\right) = \beta_0 + \beta_1 X_t + \beta_2 \texttt{post TV}_{it} + \beta_3 (X_t \times \texttt{post TV}_{it}) + \varepsilon_{it} \tag{2}$$

In the above specification, X_t represents time, in months, *post TV* is a variable that equals 1 if the focal month occurs on or after that cohort's shows are airing on television (i.e. the focal month is on or after the start of the generation cohort). For the Pre-Twitter and July 2006-2007 cohort, *post TV* takes on a value of 1 for all periods as all months in the sample are after shows in those cohorts begin airing. For the remaining cohorts, *post TV* takes on a value of zero prior January of the respective year, and a value of one thereafter. The above specification effectively allows for two diffusion curves to be fit to each generation cohort year—one pre-airing and one post-airing. I estimate the above equation separately for each of the eight generation cohorts.[3]

The primary assumption underlying the identification of Equation 2 is that the timing of a television show's premiere is not a function of the decision of the show to adopt twitter. In other words, a show's belonging to a generation cohort is exogenous to its decision to adopt Twitter.

6 Results

6.1 Visual Evidence

Figure 2 presents the diffusion of Twitter among all shows that aired at least partly after the launch of the social media service. Inspection of the graph shows an initial flat rate of growth followed by increasing rates of adoption over time. However,

[3]A model including all the data would involve random effects for the coefficient and slope, both prior to and after the airing of the shows in the focal cohort. That model is described in Equation 3 and results—presented in Column 1 of Table 6—are substantively similar to separate cohort regressions. Results by generation cohort and presented for expositional clarity.

there is no diminishing rate of cumulative adoption by the end of the period, which is characteristic of the logistic curve. Adoption appears to continue linearly after the initial period of slow adoption.

A clearer picture of diffusion is shown in Figure 3, which depicts a separate diffusion process for each generation cohort.[4] Figure 3 reveals a few trends, some of which were already evident in the summary data from Table 1. First, maximum adoption levels of all cohorts that began after the launch of Twitter range from 0.35 to 0.47. Second, pre-airing adoption occurs increasingly earlier. Third, levels of adoption at the time shows begin airing increase for each cohort.

In addition to corroborating observations from the summary statistics, the figure provides initial evidence supporting hypotheses 1 and 2. Diffusion by generation cohort more readily follows the canonical s-curve, particularly for recent cohorts. Initial levels of adoption at the beginning of the generation cohorts appear to increase over time, partly due to more prevalent pre-airing adoption. Moreover, the rate of diffusion over the period after shows begin to premiere (the post-airing period) appears to increase over time.

[4] An accompanying figure is included in Appendix Figure A1, which only includes the diffusion curve for the period after which shows began airing each respective generation cohort (i.e. the post-airing period). The figure readily shows the observations made here, with increasingly higher levels of initial adoption and faster rates of adoption across cohorts to more rapidly reach the equilibrium diffusion curve.

6.2 Parameter Estimates

Results from OLS estimation of Equation 2 for each of the eight generation cohorts are presented in Table 4. Robust standard errors are reported in parenthesis.

In order to improve interpretation and comparability across results, estimates of generation cohort shift and steepness based on the results from Table 4 are provided in Table 5. Estimates are provided separately for the pre-airing (i.e. `post TV` $= 0$) and post-airing (i.e. `post TV` $= 1$) periods. For the first four cohorts from Pre-Twitter through 2009, there were no adoptions prior to the start of the generation cohort, so the main effect of *time* in Table 4 reflects the post-airing period (this is also reflected in Table 5, where there are no pre-airing estimates for those cohorts).

For the purposes of discussing the results, I focus on Table 5. Shift is assessed as the period in which the diffusion share reached 25% penetration within the cohort. For the pre-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25)-\beta_0)/\beta_1 - M_i$, where M_i is the system stored value of the start month of the generation cohort. For the post-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25)-(\beta_0+\beta_2))/(\beta_1+\beta_3) - M_i$. Steepness is effectively the coefficient on time and is $\texttt{Steepness}_i = \beta_1$ and $\texttt{Steepness}_i = \beta_1 + \beta_3$ for the pre- and post-airing periods, respectively.

Focusing on the post-airing period, results show some corroboration of both hypotheses.

First, the first two cohorts do not experience a large change in the 25% diffusion estimate. Starting with the Jul-2007 cohort, the period at which 25% diffusion occurred consistently decreased, taking approximately four years ($47.313, p < 0.001$) for that cohort, but less than one month ($0.795, p < 0.001$) for the 2013 cohort. Second, the first five cohorts appear to have similar steepness levels. Beginning in 2010, the steepness of diffusion increased each year, from 0.100 ($p < 0.001$) in 2010 to 0.433 ($p < 0.001$) in 2013. The lag in time before the shift and steepness start to accelerate across generation cohorts provide some indication that more recent cohorts more readily incorporated the experience of the existing population and subsequent increases in both shift and steepness support hypotheses 1 and 2 during those periods.

The estimates from Table 4 are graphed by generation cohort in Figure 4. The graphs demonstrate the validity of the model when applied to the generation cohorts, particularly for later cohorts. The evidence from the raw data in Figure 3, annual estimates in Table 5, and graphs in Figure 4 provide corroborating evidence supporting both hypotheses.

6.3 Existing Population Effect

The results thus far have tested whether the shift and steepness of diffusion tends to increase over generation cohorts as the technology becomes more widely diffused in the existing population. In this section, I explicitly test whether the diffusion in the existing population affects the next generation cohort. To do so, a mixed linear model with the following specification is estimated:

$$\ln\left(\frac{s_{it}}{S_i - s_{it}}\right) = \delta_0+\delta_1 X_t+\delta_2\texttt{post TV}_{it}+\delta_3(X_t\times\texttt{post TV}_{it})+\delta_4 X_{it}+\delta_5(X_{it}\times\texttt{post TV}_{it})+\nu_{it} \tag{3}$$

where:

$$\nu_{it} = \theta_{0i} + \theta_{1i}\texttt{post TV}_{it} + \theta_{2i}X_t + \theta_{3i}(X_t \times \texttt{post TV}_{it}) + \epsilon_{it} \tag{4}$$

The error term ν_{it}, includes generation cohort errors on the shift and steepness, for periods prior to the shows airing and after, as well as the idiosyncratic error ϵ_{it}. The above specification allows for the single cohort model to be extended in a model that includes all cohorts and allows for varying shifts and steepness for each generation cohort pre- and post-airing.

The prominent feature of the model is the experience of the existing population affects the diffusion of subsequent cohorts. The variable X_{it} is a measure of diffusion extent obtained by cohort i from the existing population of prior cohorts 1 to $i-1$. Three logged measures are constructed to test the relationship. First, *prior summation* is a summation of adoption levels of all prior periods for all prior generation cohorts. If $A_{\iota\tau}$ is defined as the cumulative number of adopters for cohort ι at time τ, then for cohort, i, at time t, *prior summation* is calculated as:

$$\ln\left(\sum_{\tau=1}^{t-1}\sum_{\iota=1}^{i-1} A_{\iota\tau}\right)$$

Second, *prior level* is a summation of prior period adoption levels for all prior generation cohorts, and is calculated as:

$$\ln\left(\sum_{\iota=1}^{i-1} A_{\iota,t-1}\right)$$

Third, *prior rate* is the rate of prior period adoption for all prior generation cohorts, and is calculated as:

$$\ln\left(\frac{\sum_{\iota=1}^{i-1} A_{\iota,t-1}}{\sum_{\iota=1}^{i-1} S_{\iota,t-1}}\right)$$

where $S_{\iota,t-1}$ is the cumulative number of shows aired in cohort ι at time $t-1$.

One limitation contained in Equation 3 is that the logistic functional form is imposed on the model. As a result, the impact of prior diffusion on the shift and steepness of the curve cannot be separately identified. The coefficients δ_4 and δ_5 thus indicate the relationship between the variables of interest and generation cohort diffusion.

I estimate a mixed linear model from Equation 3 and include one of three measures of prior diffusion, *prior summation*, *prior level*, and *prior rate*. Results are reported in Table 6. Standard errors are clustered at the generation cohort level and no structure is imposed on the covariance matrix for the random effects. Column 1 excludes any prior diffusion variables, and serves to corroborate the main results from Table 4. Columns 2 through 4 exclude the Pre-Twitter cohort, as there was no prior diffusion, thus resulting in fewer observations than the baseline regression in Column 1.

The main result from these regressions is that each of the three measures of prior diffusion positively affect the rate of diffusion, in both the pre- and, more importantly, post-airing period. Focusing on the post-airing period, the estimate for the relationship between share of diffusion and *prior summation*, *prior level*, and *prior rate* is positive and statistically significant (calculated as $\delta_4 + \delta_5$ with the respective estimates being 0.18, $p < 0.001$; 0.37, $p < 0.001$; and 0.39, $p < 0.001$).

7 Robustness Tests

I identify three underlying assumptions that threaten the validity of the findings in the paper. First, show characteristics do not differentially affect the decision and the timing of adoption. Second, the sample selection criteria may bias results. Third, while the delineation of periods separating cohorts ought to be based on an the population and its underlying patterns of growth, the distinctions between cohorts can possess a degree of arbitrariness. I consider each of these assumptions in turn.

7.1 Show Characteristics

In the results discussed thus far, I have not considered how the characteristics of a show might differentially relate to its propensity to adopt. With shows airing on networks, it turns out to be the case that some networks impose more central coordination on certain activities, including social media adoption. Coordination typically results in one of two outcomes: complete adoption or complete avoidance of adoption for all shows on the network.

To test whether or not shows on network that exercise some degree of central coordination, I reproduce the results in Tables 5 and 6 using samples that exclude centrally coordinated shows. I use two definitions to assess whether a show is on a centrally coordinated network. First, the stricter definition of central coordination excludes shows if all or none of the shows on the focal show's network in the sample adopt Twitter. Second, I loosen the definition to exclude shows on networks where more than 90% or less than 10% of shows adopted. Those results are presented in Appendix Tables A1 and A2 for the strict definition and Tables A3 and A4 for the loose definition. Results from both definitions of central coordination corroborate the results discussed above.

7.2 Sample Selection

The sample is conditioned on shows that continued to air through the conclusion of the sample period. Thus, shows that were cancelled during the sample period were excluded from the diffusion. The exclusion of those shows skews the shows included in the sample towards more recent generation cohorts, as there was less time for the show to be cancelled. It may be the case that the Twitter adoption behaviors of shows that were cancelled differ from the remaining sample, thus threatening the observed results. If cancelled shows demonstrate different patterns of diffusion, the diffusion estimates may be biased.

Appendix Figure A2 shows the raw diffusion curves of each generation cohort when cancelled shows are added to the sample. What is evident is that the July 2006-2007 and 2008 cohorts have a much lower maximum adoption rate. Shows that aired in those cohorts and were cancelled were not in the sample for most of the sample period during which they would have considered adoption.

Similar to the prior section, I present in Appendix Tables A5 and A6 the annual estimates of midpoint and steepness and the random effects models of social influence, respectively. Results corroborate the primary results, implying that the exclusion of cancelled shows does not affect the results.

7.3 Definition of Generation Cohort

How generation cohorts are separated contains some ambiguity. On the one hand, the periods ought to be separated by the underlying patterns of growth in the population. On the other hand, periods of time may be arbitrary or socially constructed, or data within each period may be too sparse to observe diffusion. Ultimately, the lines separating generation cohorts must take into account the availability and density of data as well as a reasonable assessment of divisions among periods.

In the main results in Section 6, a television generation cohort included all shows premiering in a given year (for the years 2008 to 2013). Generation cohorts in the television industry might be distinguished on a semi-annual basis, since television show premieres tend to cluster around the Fall and Spring seasons. Results separating generation cohorts into half-year intervals are shown in Figures A3 and A4, with annual steepness and shift estimates in Table A7. The results are largely consistent with the annual generation cohort results. Initial levels of adoption and steepness increase, particularly for recent generation cohorts.

The results testing social influence are shown in Table A8 and are directionally the same, but not statistically significant under this definition of generation cohort. Looking at the estimates relating social influence to diffusion in the post-airing period, the estimate for *prior summation* (0.01, $p < 0.973$), *prior level* (0.20, $p < 0.234$), and *prior rate* (0.23, $p < 0.243$) are positive, but not statistically significant.

8 Discussion

For firms that are producing innovations or technologies that diffuse through populations with new generation cohorts, there are some strategic considerations to note. First, the innovation curve renews with each subsequent cohort, and the behavior of exising population affects the shift and steepness of the diffusion.

Second, new generation cohorts represent an opportunity for entrants with competing technologies to win portions of the market, even when an incumbent's technology has been widely diffused and that technology carries with it network effects. By differentiating a competing technology to appeal to a new generation, an entrant may successfully compete or replace an existing, widely diffused technology. Thinking about differentiation has typically related to product characteristics, demographic differentiation, differences in distribution channels, but it could also be that products that serve a new generational cohort could initiate growth within that cohort and then diffuse into the larger population. The generational cohort threat may partly explain why Facebook, for example, continuously opened its platform to increasingly younger people, as it is now available to thirteen year olds. It also may provide some insight as to how some online wealth management services are seeing significant levels of adoption among younger investors.

Returning to the graphical model in Figure 1, the results from Twitter diffusion into the television show generation cohorts indicate a transition in the diffusion pattern. Twitter diffused in very much the same manner for the second cohort, July 2006-2007, as it did for the initial cohort, Pre-Twitter, as depicted by Panel A. The following three cohorts from 2008 to 2010 experienced increasingly earlier shifts without substantial changes in steepness, as in panel C. The final three cohorts from 2011 to 2013 resembled Panel F, in that both the shift and steepness of the curve accelerated.

Given the transition in the nature of the diffusion curve over the generation cohorts, there is some evidence that the state into which each cohort enters the population is different, resulting in the separate paths of diffusion. At the same time, each succeeding generation cohort converges to the equilibrium diffusion path, providing some evidence that the members are part of the same population. Given both these observations, generation cohorts look to be both part of the same as the total population and part of a different set of adopters through which the innovation must diffuse.

One of the most distinctive results is that diffusion never jumps to the equilibrium path. By 2013, Twitter was the *de facto* standard for the television industry. Awareness of Twitter by that year was arguably complete, and costs to adopt (though not necessarily to manage) were effectively zero. Still, an observable process of diffusion took place, and the final generation cohort required time to reach the equilibrium path.

References

Abrahamson, E. and Rosenkopf, L. (1997). Social network effects on the extent of innovation diffusion: A computer simulation. *Organization Science*, 8(3):289–309.

Adner, R. and Kapoor, R. (2015). Innovation ecosystems and the pace of substitution: Re-examining technology s-curves. *Strategic Management Journal*, 37(4):625–648.

Alcacer, J. and Chung, W. (2007). Location strategies and knolwedge spillovers. *Management Science*, 53(5):760–776.

Anderson, P. and Tushman, M. L. (1990). Technological discontinuities and dominant designs: A cyclical model of technological change. *Administrative Science Quarterly*, 35:604–633.

Bass, F. M. (1969). A new product growth for model consumer durables. *Management Science*, 15(5):215–227.

Bresnahan, T. F. and Yin, P.-L. (2005). Economic and technical drivers of technology choice: Browsers. *Annals of Economics and Statistics / Annales d'Économie et de Statistique*, (79/80):629–670.

Caselli, F. and Coleman, W. J. (2001). Cross-country technology diffusion: The case of computers. *American Economic Review*, 91(2):328–335.

Comin, D. and Hobijn, B. (2004). Cross-country technology adoption: making the theories face the facts. *Journal of Monetary Economics*, 51(1):39–83.

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059.

Comin, D., Hobijn, B., and Rovito, E. (2006). Five facts you need to know about technology diffusion. Technical report, National Bureau of Economic Research.

Easingwood, C. J., Mahajan, V., and Muller, E. (1983). A nonuniform influence innovation diffusion model of new product acceptance. *Marketing Science*, 2(3):273–295.

Ellison, G. and Fudenberg, D. (1993). Rules of thumb for social learning. *Journal of Political Economy*, 101(4):612–643.

Farrell, J. and Saloner, G. (1985). Standardization, compatibility, and innovation. *The RAND Journal of Economics*, 16(1):70.

Farrell, J. and Saloner, G. (1986). Installed base and compatibility: Innovation, product preannouncements, and predation. *The American economic review*, pages 940–955.

Geroski, P. A. (2000). Models of technology diffusion. *Research Policy*, 29(4-5):603–625.

Goel, S., Anderson, A., Hofman, J., and Watts, D. J. (2015). The structural virality of online diffusion. *Management Science*, page 150722112809007.

Gort, M. and Klepper, S. (1982). Time paths in the diffusion of product innovations. *The Economic Journal*, 92(367):630–653.

Grajek, M. and Kretschmer, T. (2009). Usage and diffusion of cellular telephony, 19982004. *International Journal of Industrial Organization*, 27(2):238–249.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica, Journal of the Econometric Society*, 25(4):501–522.

Hannan, M. T. and Freeman, J. (1984). Structural inertia and organizational change. *American sociological review*, 49(2):149–164.

Katz, M. L. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *The American economic review*, 75(3):424–440.

Majumdar, S. K. and Venkataraman, S. (1998). Network effects and the adoption of new technology: evidence from the u.s. telecommunications industry. *Strat. Mgmt. J.*, 19(11):1045–1062.

Mansfield, E. (1963). The speed of response of firms to new techniques. *Quarterly Journal of Economics*, 77(2):290–311.

Nelson, R. R. and Winter, S. G. (1982). *An evolutionary theory of economic change.* Belknap Press.

Norton, J. A. and Bass, F. M. (1987). A diffusion theory model of adoption and substitution for successive generations of high-technology products. *Management science*, 33(9):1069–1086.

Rogers, E. M. (1995). *Diffusion of innovations.* Free Press, New York.

Rohlfs, J. H. and Hal, R. (2003). *Bandwagon effects in high-technology industries*. The MIT Press.

Saloner, G. and Shepard, A. (1995). Adoption of technologies with network effects: an empirical examination of the adoption of automated teller machines. *RAND Journal of Economics (RAND Journal of Economics)*, pages 479–501.

Tucker, C. (2008). Identifying formal and informal influence in technology adoption with network externalities. *Management Science*, 54(12):287–304.

Young, H. P. (2009). Innovation diffusion in heterogeneous populations: Contagions, social influence, and social learning. *The American Economic Review*, 99(5):1899–1924.

Table 1: Summary Statistics of generation Cohorts

Premiere Cohort	Twitter Accounts	Shows	Maximum Adoption	First Adoption Lag (months)	Initial Diffusion
Pre-Twitter	20	34	0.588	26	0.000
July 2006-2007	18	51	0.353	21	0.000
2008	16	37	0.432	3	0.000
2009	39	86	0.453	1	0.000
2010	57	149	0.383	-12	0.087
2011	105	224	0.469	-31	0.116
2012	161	412	0.391	-45	0.138
2013	236	640	0.369	-50	0.188

Note: Table presents summary statistics for each generation cohort. The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year. Twitter Accounts and Shows are the total number of accounts adopted and shows that premiered for each cohort. First Adoption Lag represents the number of months after (before if value is negative) the start of the generation cohort that the first Twitter account was adopted. Initial Diffusion is the rate of Twitter diffusion among shows at the start of the generation cohort. For the Pre-Twitter and July 2006-2007 cohorts, the start of the generation cohort is July 2006, and for each subsequent cohort it is January of the appropriate year.

Table 2: Sample Summary Statistics

	n	Mean	S.D.	Median	Min	Max
share transform	519	0.33	2.47	0.43	-5.09	4.77
share	519	0.24	0.17	0.27	0.00	0.59
time	519	614.94	18.95	615.00	579.00	647.00
post TV	519	0.73	0.44	1.00	0.00	1.00
prior summation	447	6.45	1.87	6.76	0.00	9.47
prior level	447	3.77	1.28	3.89	0.00	6.03
prior rate	447	-1.28	0.76	-1.01	-4.44	-0.53

Note: An observation is a generation cohort-month. The variable *time* is measured in system time (as reference, the launch month of Twitter in July 2006 has the value 558.)

Table 3: Correlation Table

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) share transform	1.00						
(2) share	0.95	1.00					
(3) time	0.87	0.81	1.00				
(4) post TV	0.75	0.78	0.45	1.00			
(5) prior summation	0.58	0.54	0.84	0.16	1.00		
(6) prior level	0.31	0.27	0.63	-0.08	0.94	1.00	
(7) prior rate	0.62	0.55	0.73	0.22	0.83	0.73	1.00

Note: An observation is a generation cohort-month.

Table 4: Twitter Diffusion by Premiere Year

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
time	0.094***	0.111***	0.104***	0.108***	0.199***	0.108***	0.091***	0.060***
	(0.005)	(0.003)	(0.004)	(0.005)	(0.038)	(0.004)	(0.006)	(0.004)
post TV					59.682*	-27.477***	-73.661***	-236.685***
					(22.583)	(3.966)	(4.478)	(13.409)
time x post TV					-0.099*	0.046***	0.118***	0.373***
					(0.038)	(0.006)	(0.007)	(0.021)
Constant	-56.739***	-67.078***	-62.858***	-65.027***	-120.809***	-67.519***	-56.893***	-38.831***
	(3.382)	(1.593)	(2.174)	(2.949)	(22.402)	(2.611)	(3.433)	(2.403)
Model Year	Pre-Twitter	Jul 2006-07	2008	2009	2010	2011	2012	2013
Calendar Months	64	69	69	59	60	67	69	62
F-Stat	293.1	1827.1	859.3	515.4	394.0	4668.1	1998.6	1606.0
Adj R-squared	0.89	0.97	0.93	0.92	0.97	0.99	0.98	0.98

* $p < 0.05$; ** $p < 0.01$; *** < 0.001.

Note: Table presents results from separate OLS models for each generation cohort. The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year. The dependent variable is a function of the share of shows premiering in year i that adopted Twitter at time t and the maximum share for that premiere year, $\ln(s_{it}/(S_i - s_{it}))$. Robust standard errors are reported in parentheses. For the Pre-Twitter, July 2006-2007, 2008, and 2009 cohorts, the coefficient on the `time` variable is for the post airing period, as there was no diffusion of Twitter prior to the start of the generation cohort.

Table 5: Estimates for Diffusion and Steepness by Year

Premiere Cohort	Maximum Adoption	Pre-airing 25% Diffusion	SE	Pre-airing Steepness	SE	Post-airing 25% Diffusion	SE	Post-airing Steepness	SE
Pre-Twitter	0.588	na	na	na	na	48.316	0.983	0.094	0.005
July 2006-2007	0.353	na	na	na	na	47.313	0.451	0.111	0.003
2008	0.432	na	na	na	na	29.370	0.680	0.104	0.004
2009	0.453	na	na	na	na	18.904	0.906	0.108	0.005
2010	0.383	7.123	2.438	0.199	0.040	13.459	0.414	0.100	0.005
2011	0.469	16.312	1.487	0.108	0.004	5.703	0.534	0.154	0.005
2012	0.391	7.460	1.600	0.091	0.006	3.614	0.236	0.208	0.005
2013	0.369	17.443	2.581	0.060	0.004	0.795	0.138	0.433	0.022

Note: Estimates for diffusion and steepness provided by generation cohort and based on estimation of Equation 2 for each generation cohort (results presented in Table 4). The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year. Both estimates are provided for the pre-airing and post-airing periods of each generation cohort. The post-airing period begins after shows in that generation cohort begin airing. For the Pre-Twitter and July 2006-2007 cohorts, the post-airing period begins July 2006, and for each subsequent cohort it is January of the appropriate year. For each generation cohort until 2009, there is no pre-airing estimate, as all adoptions took place in the post-airing period. Diffusion curves are characterized by maximum adoption, shift, and steepness. Maximum adoption, S_i is the highest share of shows of generation cohort i that adopted Twitter and is treated as exogenous in the model. Shift is displayed as the 25% Diffusion estimate, which represents the number of months after the generation cohort begins airing that Twitter diffusion for that generation cohort is 25%. For the pre-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25) - \beta_0)/\beta_1 - M_i$, where M_i is the system stored value of the start month of the generation cohort. For the post-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25) - (\beta_0 + \beta_2))/(\beta_1 + \beta_3) - M_i$ Steepness in the pre-airing period is $\texttt{Steepness}_i = \beta_1$ and in the post-airing period is $\texttt{Steepness}_i = \beta_1 + \beta_3$. Standard errors for all estimates are included.

Table 6: Random effects models of Twitter Diffusion

	(1)	(2)	(3)	(4)
time	0.078***	0.021	0.033**	0.055***
	(0.012)	(0.011)	(0.012)	(0.006)
post TV	-49.666	-83.373**	-78.909**	-66.341*
	(29.261)	(27.925)	(30.412)	(30.431)
time x post TV	0.082	0.140**	0.131**	0.109*
	(0.046)	(0.043)	(0.048)	(0.048)
prior summation		0.425***		
		(0.093)		
prior summation x post TV		-0.246*		
		(0.101)		
prior level			0.500***	
			(0.134)	
prior level x post TV			-0.133	
			(0.145)	
prior rate				0.493***
				(0.133)
prior rate x post TV				-0.105
				(0.139)
Constant	-49.440***	-17.616**	-24.025***	-34.571***
	(6.901)	(6.258)	(6.514)	(3.978)
var(Constant)	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
var(time x pre TV)	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
var(post TV)	4026.547***	4789.919***	4881.654***	4972.300***
	(3606.345)	(3937.339)	(3986.211)	(3759.831)
var(time x post TV)	0.010***	0.011***	0.012***	0.012***
	(0.009)	(0.009)	(0.009)	(0.009)
cov(time, time x post TV)	-5.844***	-5.847***	-5.679***	-6.065***
	(0.561)	(0.551)	(0.531)	(0.552)
var(Residual)	0.211***	0.154***	0.144***	0.144***
	(0.038)	(0.039)	(0.035)	(0.035)
Start Year-Months	519	447	447	447
Gen Cohorts (clusters)	8	7	7	7
Log Likelihood	-375.4	-256.8	-244.8	-241.7

* $p < 0.05$; ** $p < 0.01$; *** < 0.001.

Note: Table presents results from a mixed linear model (MLM) estimated using maximum likelihood. The dependent variable is a function of the share of shows premiering in year i that adopted Twitter at time t and the maximum share for that premiere year, $\ln(s_{it}/(S_i - s_{it}))$. Standard errors, reported in parentheses, are clustered at the premiere year. No structure is imposed on the covariance matrix for the random effects. Note that in Columns 2 to 4, the Pre-Twitter cohort is excluded from the sample, as there is no data available for *prior level*, *prior summation*, or *prior rate*.

Figure 1: Illustrative Examples of Generation Cohort Diffusion

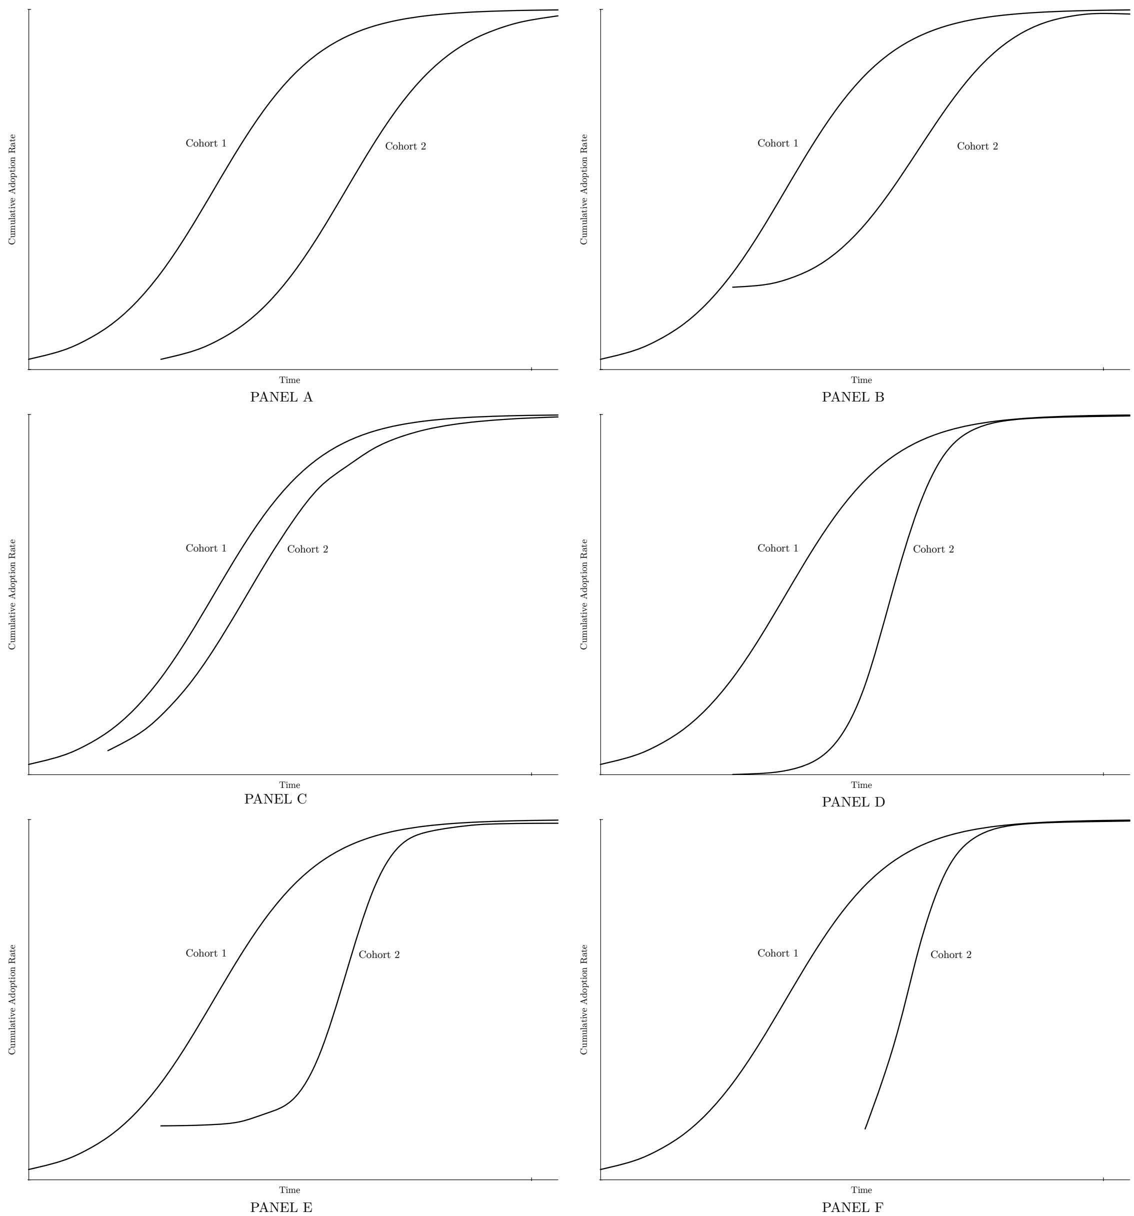

Note: Each panel depicts an illustrative example of two generation cohorts.

Figure 2: Twitter Diffusion Among Television Shows

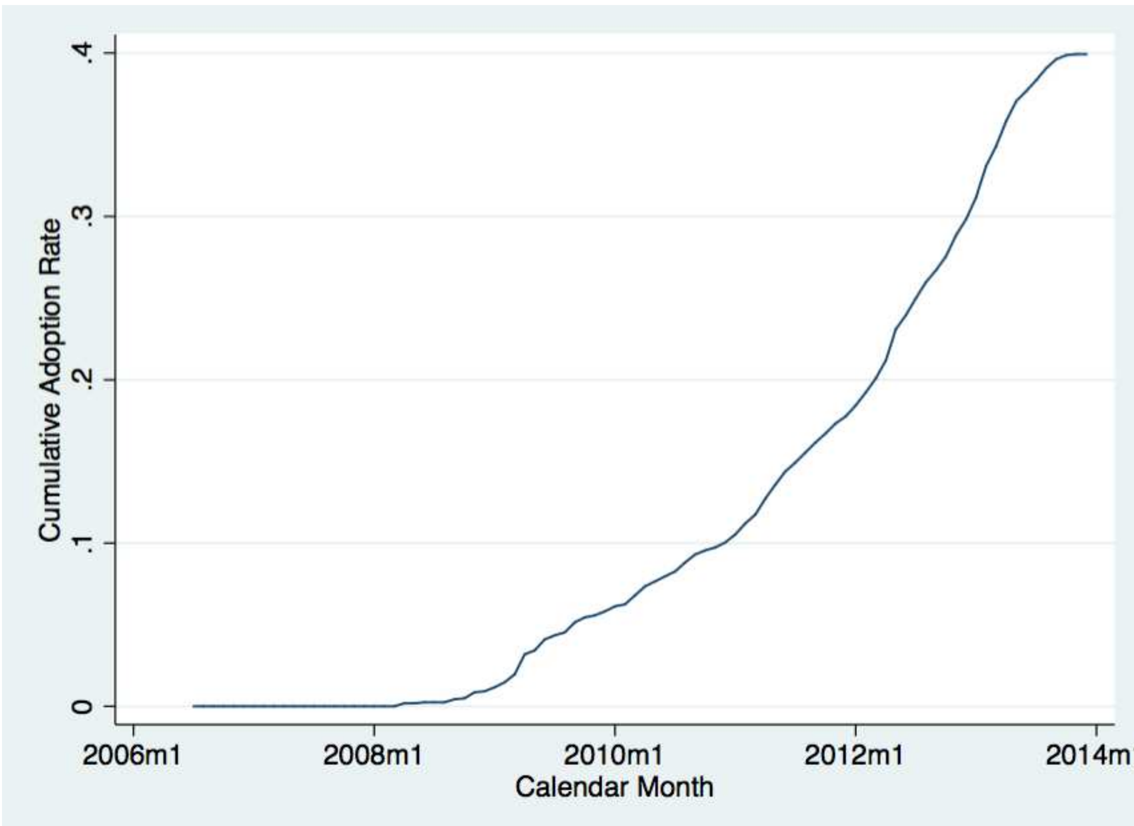

Note: The graph depicts s_{it} among all television shows and is calculated as the cumulative number of shows i that adopted Twitter at time t divided by the total shows that premiered from 2005 to 2013 and at aired least partly after the launch of Twitter in July 2006.

Figure 3: Twitter Diffusion by Premiere Year of Show

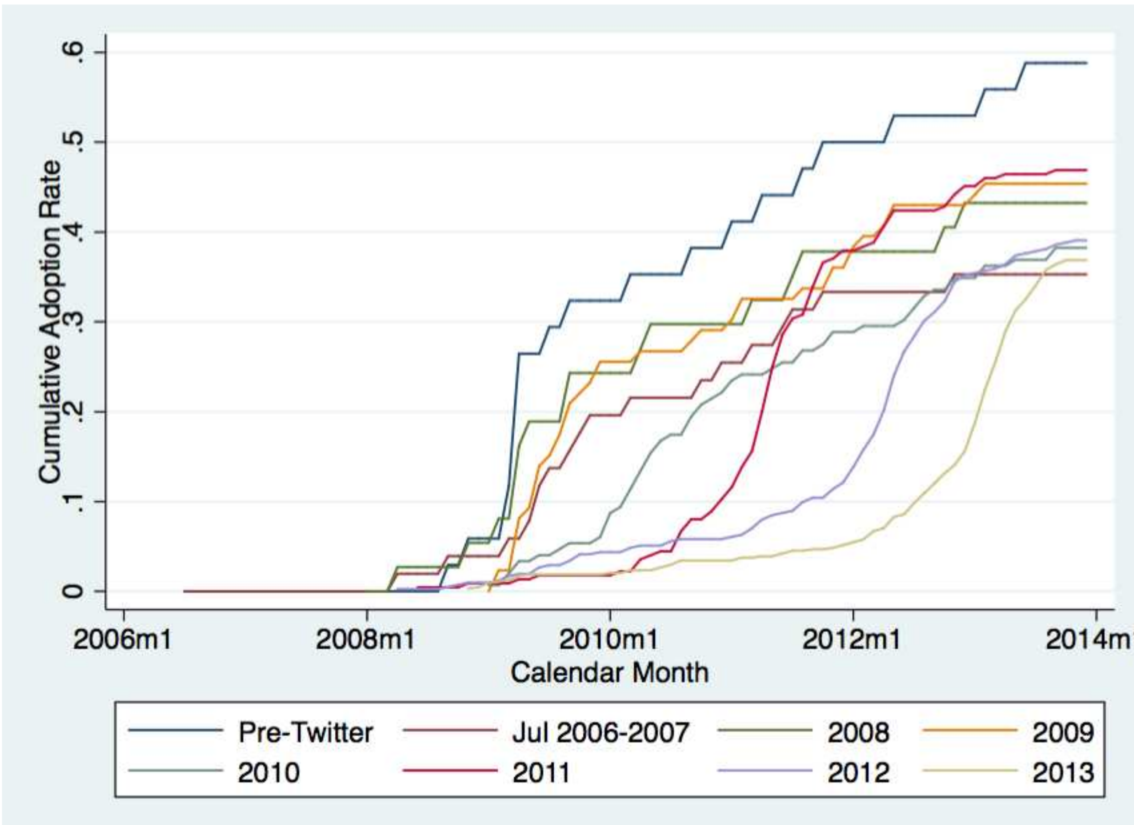

Note: Each line depicts s_{it} for a given cohort of shows that premiered that year, which is calculated as the cumulative number of shows that adopted Twitter at time t divided by the total shows in each respective cohort during the entire sample period. The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year.

Figure 4: Predicted Twitter Diffusion by generation Cohort

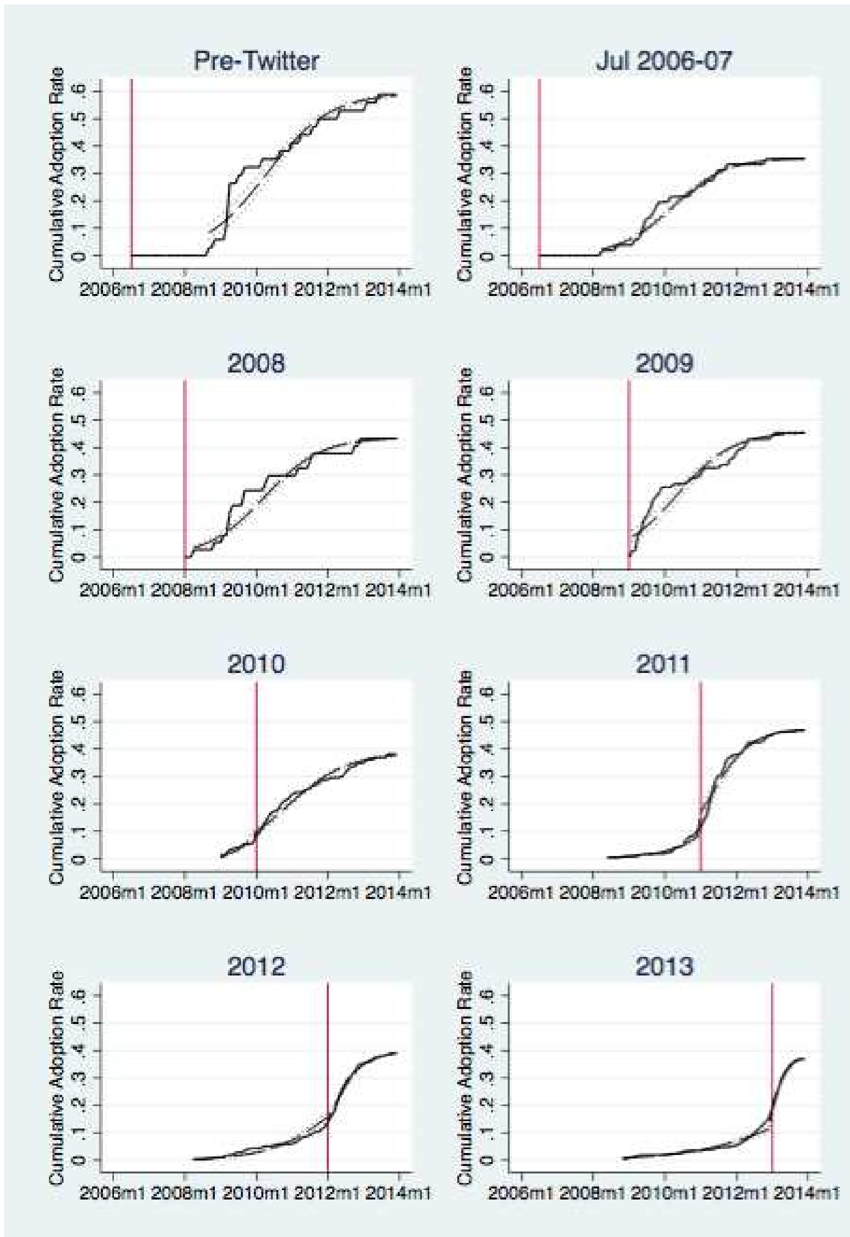

Note: Each pane shows the diffusion for a single generation cohort, compared to predicted values derived from estimates in Table 4. The solid line represents actual diffusion, the long hypenated line is predicted diffusion and the dotted lines are the 95% confidence intervals for the predicted values. The vertical red lines indicate the period at which shows began airing for that generation cohort (hence separating the pre-airing portion from the post-airing portion of the diffusion.)

Table A1: Diffusion and Steepness Estimates Excluding Centrally Coordinated Shows (Strict Definition)

Premiere Cohort	Maximum Adoption	Pre-airing 25% Diffusion	SE	Pre-airing Steepness	SE	Post-airing 25% Diffusion	SE	Post-airing Steepness	SE
Pre-Twitter	0.600	na	na	na	na	45.337	1.150	0.094	0.005
July 2006-2007	0.347	na	na	na	na	47.689	0.393	0.111	0.002
2008	0.432	na	na	na	na	29.370	0.680	0.104	0.004
2009	0.434	na	na	na	na	18.406	0.855	0.111	0.004
2010	0.371	8.345	2.985	0.182	0.041	14.054	0.417	0.103	0.004
2011	0.460	15.983	1.394	0.107	0.004	5.587	0.520	0.153	0.005
2012	0.374	7.781	1.559	0.089	0.005	3.630	0.267	0.206	0.005
2013	0.340	14.824	2.349	0.060	0.004	0.784	0.179	0.422	0.022

Note: This table replicates results for Table 5 using a sample that excludes shows on networks where all or no shows adopted Twitter. Estimates for diffusion and steepness provided by generation cohort and based on estimation of Equation 2 for each generation cohort (results presented in Table 4). The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year. Both estimates are provided for the pre-airing and post-airing periods of each generation cohort. The post-airing period begins after shows in that generation cohort begin airing. For the Pre-Twitter and July 2006-2007 cohorts, the post-airing period begins July 2006, and for each subsequent cohort it is January of the appropriate year. For each generation cohort until 2009, there is no pre-airing estimate, as all adoptions took place in the post-airing period. Diffusion curves are characterized by maximum adoption, shift, and steepness. Maximum adoption, S_i is the highest share of shows of generation cohort i that adopted Twitter and is treated as exogenous in the model. Shift is displayed as the 25% Diffusion estimate, which represents the number of months after the generation cohort begins airing that Twitter diffusion for that generation cohort is 25%. For the pre-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25)-\beta_0)/\beta_1 - M_i$, where M_i is the system stored value of the start month of the generation cohort. For the post-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25)-(\beta_0+\beta_2))/(\beta_1+\beta_3) - M_i$ Steepness in the pre-airing period is $\texttt{Steepness}_i = \beta_1$ and in the post-airing period is $\texttt{Steepness}_i = \beta_1 + \beta_3$. Standard errors for all estimates are included.

Table A2: Random effects models Excluding Centrally Coordinated Shows (Strict Definition)

	(1)	(2)	(3)	(4)
time	0.077***	0.022*	0.033**	0.055***
	(0.011)	(0.010)	(0.011)	(0.006)
post TV	-49.578	-82.712**	-78.573**	-66.400*
	(28.029)	(26.901)	(29.209)	(29.176)
time x post TV	0.082	0.139***	0.130**	0.109*
	(0.044)	(0.042)	(0.046)	(0.045)
prior summation		0.416***		
		(0.083)		
prior summation x post TV		-0.259**		
		(0.092)		
prior level			0.495***	
			(0.123)	
prior level x post TV			-0.178	
			(0.138)	
prior rate				0.491***
				(0.123)
prior rate x post TV				-0.154
				(0.130)
Constant	-48.861***	-18.180**	-24.147***	-34.598***
	(6.620)	(5.560)	(6.100)	(3.782)
var(Constant)	0.000***	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
var(time x pre TV)	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
var(post TV)	3704.279***	4385.117***	4455.412***	4541.901***
	(3340.481)	(3652.690)	(3714.230)	(3715.218)
var(time x post TV)	0.009***	0.010***	0.011***	0.011***
	(0.008)	(0.009)	(0.009)	(0.009)
cov(time, time x post TV)	-5.855***	-5.810***	-5.663***	-5.982***
	(0.586)	(0.549)	(0.529)	(0.564)
var(Residual)	0.197***	0.147***	0.140***	0.140***
	(0.036)	(0.038)	(0.035)	(0.035)
Start Year-Months	519	447	447	447
Start Years (clusters)	8	7	7	7
Log Likelihood	-357.0	-246.7	-237.2	-234.5

* $p < 0.05$; ** $p < 0.01$; *** < 0.001.

Note: This table replicates results for Table 6 using a sample that excludes shows on networks where all or no shows adopted Twitter. Table presents results from a mixed linear model (MLM) estimated using maximum likelihood. The dependent variable is a function of the share of shows premiering in year i that adopted Twitter at time t and the maximum share for that premiere year, $\ln(s_{it}/(S_i - s_{it}))$. Standard errors, reported in parentheses, are clustered at the premiere year. No structure is imposed on the covariance matrix for the random effects. Note that in Columns 2 to 4, the Pre-Twitter cohort is excluded from the sample, as there is no data available for *prior level*, *prior summation*, or *prior rate*.

Table A3: Diffusion and Steepness Estimates Excluding Centrally Coordinated Shows (Loose Definition)

Premiere Cohort	Maximum Adoption	Pre-airing 25% Diffusion	SE	Pre-airing Steepness	SE	Post-airing 25% Diffusion	SE	Post-airing Steepness	SE
Pre-Twitter	0.615	na	na	na	na	45.557	1.471	0.092	0.007
July 2006-2007	0.326	na	na	na	na	47.521	0.438	0.106	0.002
2008	0.382	na	na	na	na	28.094	0.820	0.096	0.004
2009	0.416	na	na	na	na	17.868	0.934	0.107	0.004
2010	0.405	17.263	5.635	0.111	0.030	14.198	0.418	0.103	0.005
2011	0.463	16.736	1.378	0.102	0.004	6.410	0.388	0.157	0.004
2012	0.380	8.157	1.753	0.086	0.006	3.987	0.223	0.208	0.005
2013	0.332	13.857	2.689	0.062	0.005	0.824	0.183	0.418	0.021

Note: This table replicates results for Table 5 using a sample that excludes shows on networks with total adoption rates of less than 10% or greater than 90%. Estimates for diffusion and steepness provided by generation cohort and based on estimation of Equation 2 for each generation cohort (results presented in Table 4). The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year. Both estimates are provided for the pre-airing and post-airing periods of each generation cohort. The post-airing period begins after shows in that generation cohort begin airing. For the Pre-Twitter and July 2006-2007 cohorts, the post-airing period begins July 2006, and for each subsequent cohort it is January of the appropriate year. For each generation cohort until 2009, there is no pre-airing estimate, as all adoptions took place in the post-airing period. Diffusion curves are characterized by maximum adoption, shift, and steepness. Maximum adoption, S_i is the highest share of shows of generation cohort i that adopted Twitter and is treated as exogenous in the model. Shift is displayed as the 25% Diffusion estimate, which represents the number of months after the generation cohort begins airing that Twitter diffusion for that generation cohort is 25%. For the pre-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25) - \beta_0)/\beta_1 - M_i$, where M_i is the system stored value of the start month of the generation cohort. For the post-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25) - (\beta_0 + \beta_2))/(\beta_1 + \beta_3) - M_i$ Steepness in the pre-airing period is $\texttt{Steepness}_i = \beta_1$ and in the post-airing period is $\texttt{Steepness}_i = \beta_1 + \beta_3$. Standard errors for all estimates are included.

Table A4: Random effects models Excluding Centrally Coordinated Shows (Loose Definition)

	(1)	(2)	(3)	(4)
time	0.076***	0.023*	0.035***	0.055***
	(0.009)	(0.011)	(0.011)	(0.005)
post TV	-48.825	-80.530**	-75.456*	-64.253*
	(26.846)	(27.312)	(29.589)	(29.164)
time x post TV	0.081	0.135**	0.125**	0.105*
	(0.042)	(0.043)	(0.046)	(0.045)
prior summation		0.403***		
		(0.090)		
prior summation x post TV		-0.236*		
		(0.113)		
prior level			0.460***	
			(0.124)	
prior level x post TV			-0.097	
			(0.152)	
prior rate				0.459***
				(0.124)
prior rate x post TV				-0.077
				(0.140)
Constant	-48.059***	-18.588**	-25.365***	-34.757***
	(5.587)	(6.382)	(5.941)	(3.498)
var(Constant)	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
var(time x pre TV)	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
var(post TV)	3665.797***	4378.778***	4438.846***	4561.225***
	(3267.621)	(3551.883)	(3616.474)	(3779.755)
var(time x post TV)	0.009***	0.010***	0.011***	0.011***
	(0.008)	(0.008)	(0.009)	(0.009)
cov(time, time x post TV)	-5.976***	-5.925***	-5.745***	-6.085***
	(0.587)	(0.541)	(0.517)	(0.561)
var(Residual)	0.221***	0.163***	0.154***	0.154***
	(0.051)	(0.048)	(0.042)	(0.042)
Start Year-Months	515	443	443	443
Start Years (clusters)	8	7	7	7
Log Likelihood	-382.7	-266.2	-255.4	-252.8

* $p < 0.05$; ** $p < 0.01$; *** < 0.001.

Note: This table replicates results for Table 6 using a sample that excludes shows on networks with total adoption rates of less than 10% or greater than 90%. Table presents results from a mixed linear model (MLM) estimated using maximum likelihood. The dependent variable is a function of the share of shows premiering in year i that adopted Twitter at time t and the maximum share for that premiere year, $\ln(s_{it}/(S_i - s_{it}))$. Standard errors, reported in parentheses, are clustered at the premiere year. No structure is imposed on the covariance matrix for the random effects. Note that in Columns 2 to 4, the Pre-Twitter cohort is excluded from the sample, as there is no data available for *prior level*, *prior summation*, or *prior rate*.

Table A5: Diffusion and Steepness Estimates Including Cancelled Shows

Premiere Cohort	Maximum Adoption	Pre-airing 25% Diffusion	SE	Pre-airing Steepness	SE	Post-airing 25% Diffusion	SE	Post-airing Steepness	SE
Pre-Twitter	0.256	na	na	na	na	47.673	0.509	0.096	0.002
July 2006-2007	0.136	na	na	na	na	47.868	0.726	0.094	0.003
2008	0.144	na	na	na	na	30.528	1.127	0.090	0.005
2009	0.244	10.110	2.321	0.335	0.059	13.565	1.686	0.093	0.006
2010	0.300	14.506	2.972	0.145	0.014	6.394	1.020	0.089	0.004
2011	0.373	8.765	0.683	0.136	0.003	1.190	1.088	0.133	0.007
2012	0.382	4.747	1.060	0.107	0.005	1.243	0.561	0.194	0.007
2013	0.387	16.288	2.559	0.063	0.004	0.623	0.122	0.421	0.016

Note: This table replicates results for Table 5 using a sample that includes shows that were cancelled during the sample period. Estimates for diffusion and steepness provided by generation cohort and based on estimation of Equation 2 for each generation cohort (results presented in Table 4). The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year. Both estimates are provided for the pre-airing and post-airing periods of each generation cohort. The post-airing period begins after shows in that generation cohort begin airing. For the Pre-Twitter and July 2006-2007 cohorts, the post-airing period begins July 2006, and for each subsequent cohort it is January of the appropriate year. For each generation cohort until 2009, there is no pre-airing estimate, as all adoptions took place in the post-airing period. Diffusion curves are characterized by maximum adoption, shift, and steepness. Maximum adoption, S_i is the highest share of shows of generation cohort i that adopted Twitter and is treated as exogenous in the model. Shift is displayed as the 25% Diffusion estimate, which represents the number of months after the generation cohort begins airing that Twitter diffusion for that generation cohort is 25%. For the pre-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25)-\beta_0)/\beta_1 - M_i$, where M_i is the system stored value of the start month of the generation cohort. For the post-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25)-(\beta_0+\beta_2))/(\beta_1+\beta_3) - M_i$ Steepness in the pre-airing period is $\texttt{Steepness}_i = \beta_1$ and in the post-airing period is $\texttt{Steepness}_i = \beta_1 + \beta_3$. Standard errors for all estimates are included.

Table A6: Random effects models Including Cancelled Shows

	(1)	(2)	(3)	(4)
time	0.093***	0.010***	0.035***	0.049***
	(0.017)	(0.003)	(0.009)	(0.009)
post TV	-31.572	-62.117*	-63.598	-52.461
	(32.733)	(30.764)	(33.503)	(33.999)
time x post TV	0.053	0.100*	0.101	0.088
	(0.052)	(0.047)	(0.052)	(0.053)
prior summation		0.688***		
		(0.065)		
prior summation x post TV		0.241		
		(0.129)		
prior level			0.639***	
			(0.064)	
prior level x post TV			0.522***	
			(0.103)	
prior rate				0.712***
				(0.081)
prior rate x post TV				0.454***
				(0.113)
Constant	-58.890***	-13.406***	-26.788***	-30.254***
	(10.020)	(1.484)	(5.178)	(5.298)
var(Constant)	0.000	0.000	0.000	0.133*
	(0.000)	(0.000)	(0.000)	(0.122)
var(time x pre TV)	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
var(post TV)	4043.842***	5627.339***	5791.503***	5968.406***
	(3519.996)	(4366.820)	(4987.052)	(4420.113)
var(time x post TV)	0.010***	0.013***	0.014***	0.014***
	(0.008)	(0.010)	(0.012)	(0.011)
cov(time, time x post TV)	-6.477***	-6.092***	-5.825*	-6.138***
	(0.631)	(0.490)	(2.299)	(0.731)
var(Residual)	0.250***	0.128***	0.095***	0.095***
	(0.034)	(0.022)	(0.023)	(0.022)
Start Year-Months	567	486	486	486
Start Years (clusters)	8	7	7	7
Log Likelihood	-448.8	-234.0	-166.8	-164.7

* p < 0.05; ** p < 0.01; *** < 0.001.

Note: This table replicates results for Table 6 using a sample that includes shows that were cancelled during the sample period. Table presents results from a mixed linear model (MLM) estimated using maximum likelihood. The dependent variable is a function of the share of shows premiering in year i that adopted Twitter at time t and the maximum share for that premiere year, $\ln(s_{it}/(S_i - s_{it}))$. Standard errors, reported in parentheses, are clustered at the premiere year. No structure is imposed on the covariance matrix for the random effects. Note that in Columns 2 to 4, the Pre-Twitter cohort is excluded from the sample, as there is no data available for *prior level*, *prior summation*, or *prior rate*.

Table A7: Diffusion and Steepness Estimates for Half-Year Cohorts

Premiere Cohort	Maximum Adoption	Pre-airing 25% Diffusion	SE	Pre-airing Steepness	SE	Post-airing 25% Diffusion	SE	Post-airing Steepness	SE
Pre-Twitter	0.588	na	na	na	na	48.316	0.983	0.094	0.005
July 2006-2007	0.353	na	na	na	na	47.313	0.451	0.111	0.003
2008h1	0.250	na	na	na	na	34.873	1.668	0.075	0.009
2008h2	0.571	na	na	na	na	17.706	0.679	0.125	0.005
2009h1	0.448	na	na	na	na	11.194	1.207	0.099	0.004
2009h2	0.456	2.646	1.955	0.524	0.214	14.606	0.631	0.108	0.004
2010h1	0.371	11.178	5.295	0.129	0.046	9.686	0.650	0.119	0.004
2010h2	0.391	7.149	1.369	0.144	0.011	9.308	0.772	0.095	0.008
2011h1	0.465	5.798	1.435	0.182	0.016	3.160	0.434	0.144	0.005
2011h2	0.472	17.898	4.357	0.081	0.008	-1.079	0.557	0.143	0.005
2012h1	0.364	6.059	1.008	0.088	0.003	0.485	0.839	0.162	0.011
2012h2	0.409	7.758	1.857	0.073	0.004	-1.681	0.231	0.229	0.006
2013h1	0.347	9.873	2.495	0.066	0.005	-2.776	0.731	0.306	0.024
2013h2	0.394	9.494	3.840	0.067	0.006	-5.228	3.054	0.450	0.164

Note: This table replicates results for Table 5 using a sample that defines a generation cohort in half-year intervals from 2008 to 2013. Estimates for diffusion and steepness provided by generation cohort and based on estimation of Equation 2 for each generation cohort (results presented in Table 4). The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that six month calendar period. Both estimates are provided for the pre-airing and post-airing periods of each generation cohort. The post-airing period begins after shows in that generation cohort begin airing. For the Pre-Twitter and July 2006-2007 cohorts, the post-airing period begins July 2006, and for each subsequent cohort it is January or July of the appropriate six month period. For each generation cohort until 2009, there is no pre-airing estimate, as all adoptions took place in the post-airing period. Diffusion curves are characterized by maximum adoption, shift, and steepness. Maximum adoption, S_i is the highest share of shows of generation cohort i that adopted Twitter and is treated as exogenous in the model. Shift is displayed as the 25% Diffusion estimate, which represents the number of months after the generation cohort begins airing that Twitter diffusion for that generation cohort is 25%. For the pre-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25) - \beta_0)/\beta_1 - M_i$, where M_i is the system stored value of the start month of the generation cohort. For the post-airing period, it is calculated as $\texttt{25\% Diffusion}_i = (-ln(1-0.25) - (\beta_0 + \beta_2))/(\beta_1 + \beta_3) - M_i$ Steepness in the pre-airing period is $\texttt{Steepness}_i = \beta_1$ and in the post-airing period is $\texttt{Steepness}_i = \beta_1 + \beta_3$. Standard errors for all estimates are included.

Table A8: Random effects models for Half-Year Cohorts

	(1)	(2)	(3)	(4)
time	0.075***	0.058***	0.060***	0.069***
	(0.005)	(0.009)	(0.009)	(0.005)
post TV	-33.490***	-45.564***	-44.400***	-38.459***
	(9.178)	(11.403)	(11.639)	(10.674)
time x post TV	0.057***	0.078***	0.074***	0.064***
	(0.015)	(0.020)	(0.019)	(0.017)
prior summation		0.144		
		(0.082)		
prior summation x post TV		-0.137		
		(0.239)		
prior level			0.198	
			(0.112)	
prior level x post TV			0.000	
			(0.207)	
prior rate				0.192
				(0.111)
prior rate x post TV				0.037
				(0.234)
Constant	-47.319***	-37.901***	-38.692***	-43.047***
	(2.662)	(4.784)	(5.049)	(3.023)
var(Constant)	0.290	0.394	0.436	0.367
	(0.225)	(0.292)	(0.305)	(0.289)
var(time x pre TV)	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
var(post TV)	705.104***	804.386***	906.329***	884.274***
	(693.941)	(766.928)	(791.683)	(791.043)
var(time x post TV)	0.002***	0.002***	0.002***	0.002***
	(0.002)	(0.002)	(0.002)	(0.002)
cov(time, time x post TV)	-4.636***	-4.797***	-4.779***	-4.879***
	(0.624)	(0.666)	(0.640)	(0.665)
var(Residual)	0.325***	0.318***	0.314***	0.314***
	(0.079)	(0.087)	(0.088)	(0.088)
Start Year-Months	874	802	802	802
Start Years (clusters)	14	13	13	13
Log Likelihood	-815.0	-741.6	-737.9	-737.1

* $p < 0.05$; ** $p < 0.01$; *** < 0.001.

Note: This table replicates results for Table 6 using a sample that defines a generation cohort in half-year intervals from 2008 to 2013. Table presents results from a mixed linear model (MLM) estimated using maximum likelihood. The dependent variable is a function of the share of shows premiering in year i that adopted Twitter at time t and the maximum share for that premiere year, $\ln(s_{it}/(S_i - s_{it}))$. Standard errors, reported in parentheses, are clustered at the premiere year. No structure is imposed on the covariance matrix for the random effects. Note that in Columns 2 to 4, the Pre-Twitter cohort is excluded from the sample, as there is no data available for *prior level*, *prior summation*, or *prior rate*.

Figure A1: Twitter Diffusion by Premiere Year of Show (Post-Airing Periods Only)

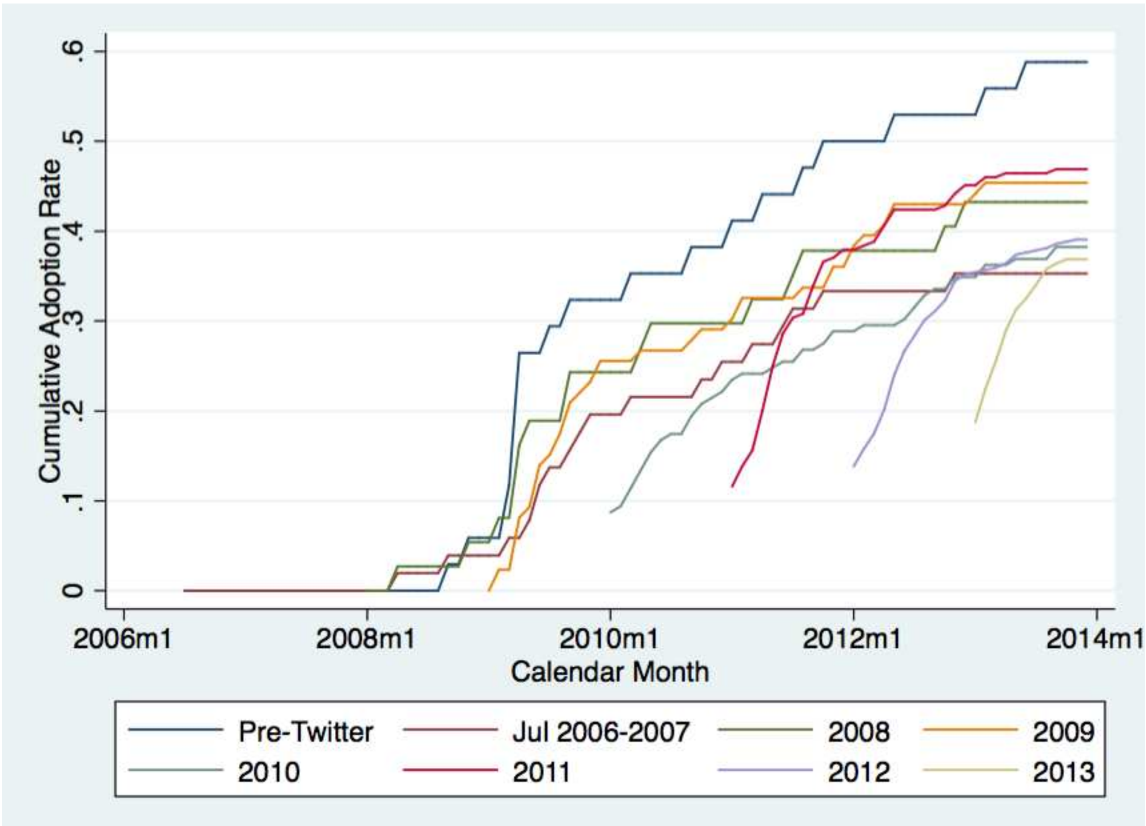

Note: This chart replicates Figure 2, except each generation cohort only shows diffusion for the period after which television shows began airing in that cohort. Each line depicts s_{it} for a given cohort of shows that premiered that year, which is calculated as the cumulative number of shows that adopted Twitter at time t divided by the total shows in each respective cohort during the entire sample period. The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year.

Figure A2: Twitter Diffusion by Premiere Year of Show (Including Cancelled Shows)

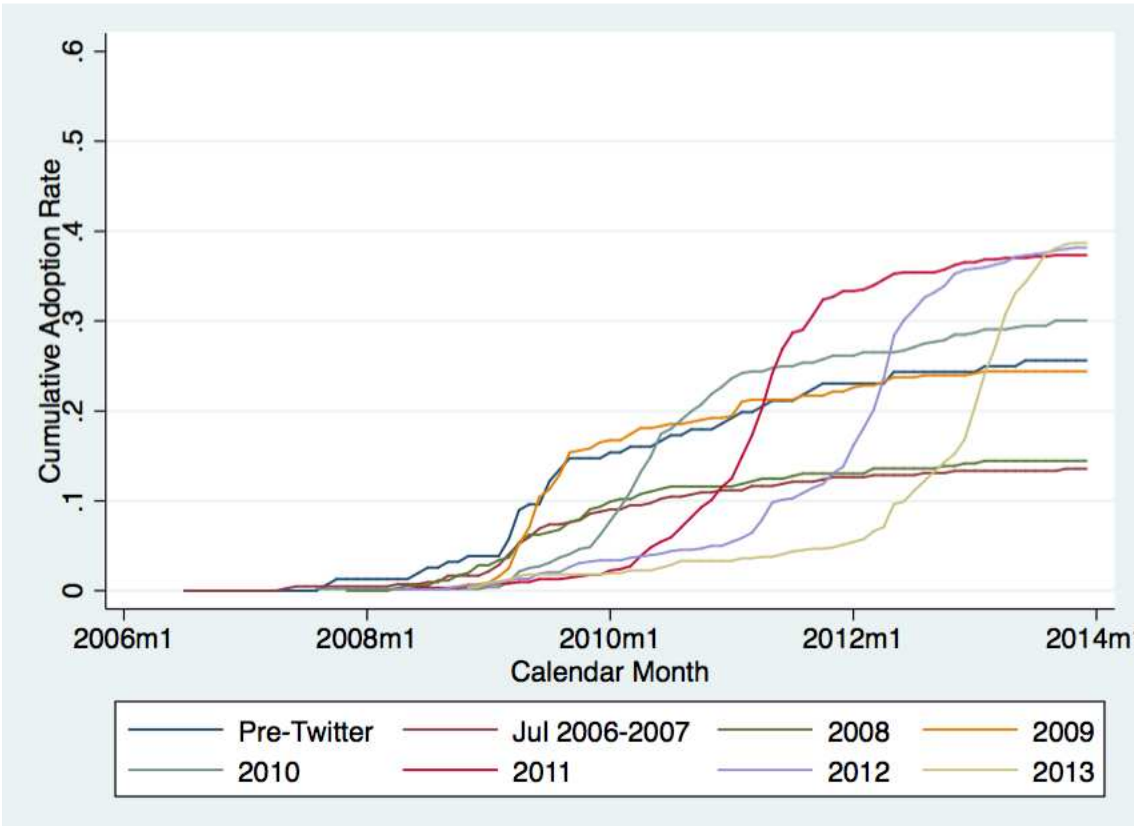

Note: This figure replicates Figure 3, but includes shows that were cancelled prior to the completion of the sample period. Each line depicts s_{it} for a given cohort of shows that premiered that year, which is calculated as the cumulative number of shows that adopted Twitter at time t divided by the total shows in each respective cohort during the entire sample period. The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort include shows that aired during that calendar year.

Figure A3: Twitter Diffusion by Semi-Annual generation Cohort (Pre-Twitter to 2010h2)

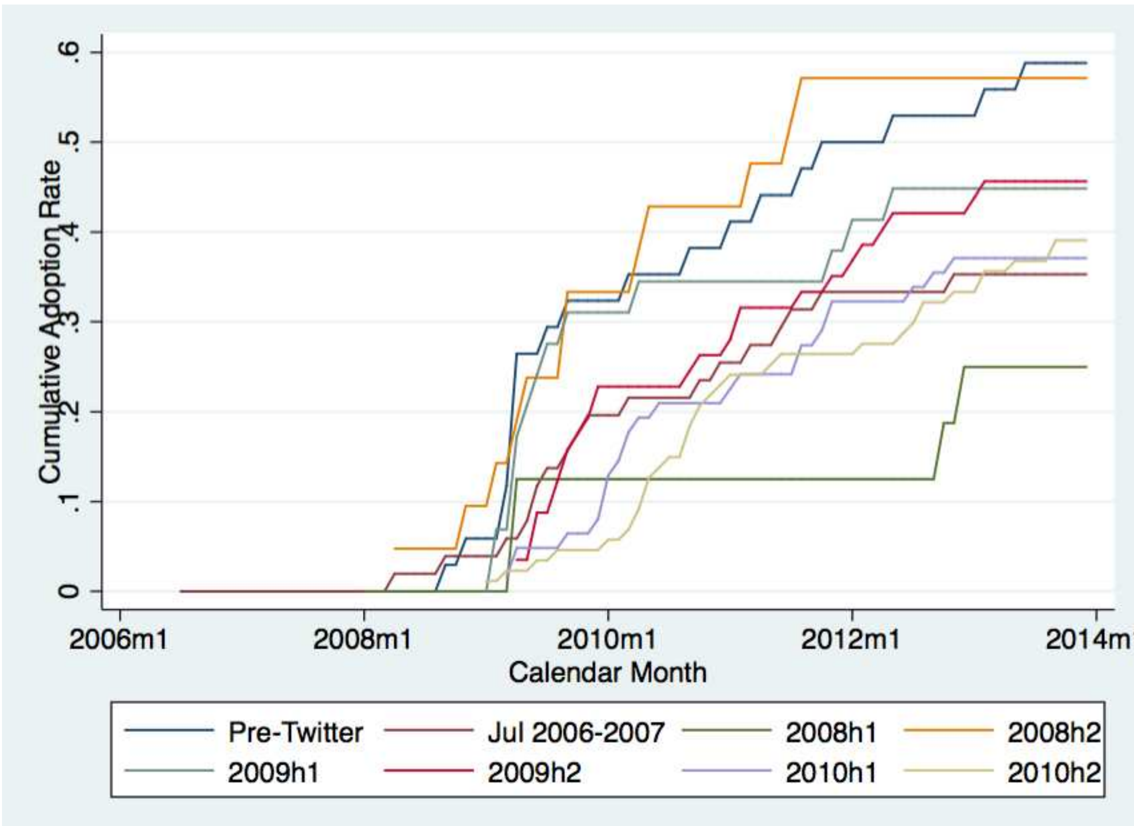

Note: This figure replicates Figure 3, but separates each generation cohort semi-annually from 2008 to 2013. For clarity, the half-year cohorts are split between Figures A3 and A4. Each line depicts s_{it} for a given cohort of shows that prehmiered that half-year, which is calculated as the cumulative number of shows that adopted Twitter at time t divided by the total shows in each respective cohort during the entire sample period. The Pre-Twitter cohort includes shows that premiered from January 2005 to July 2006 and continued airing, in part, after the launch of Twitter (July 15, 2006). The July 2006-2007 cohort includes 2006 shows that premiered after Twitter's July launch and all shows premiering in 2007. Each subsequent cohort includes shows that aired during that half-year.

Figure A4: Twitter Diffusion by Semi-Annual generation Cohort (2011h1 to 2013h2)

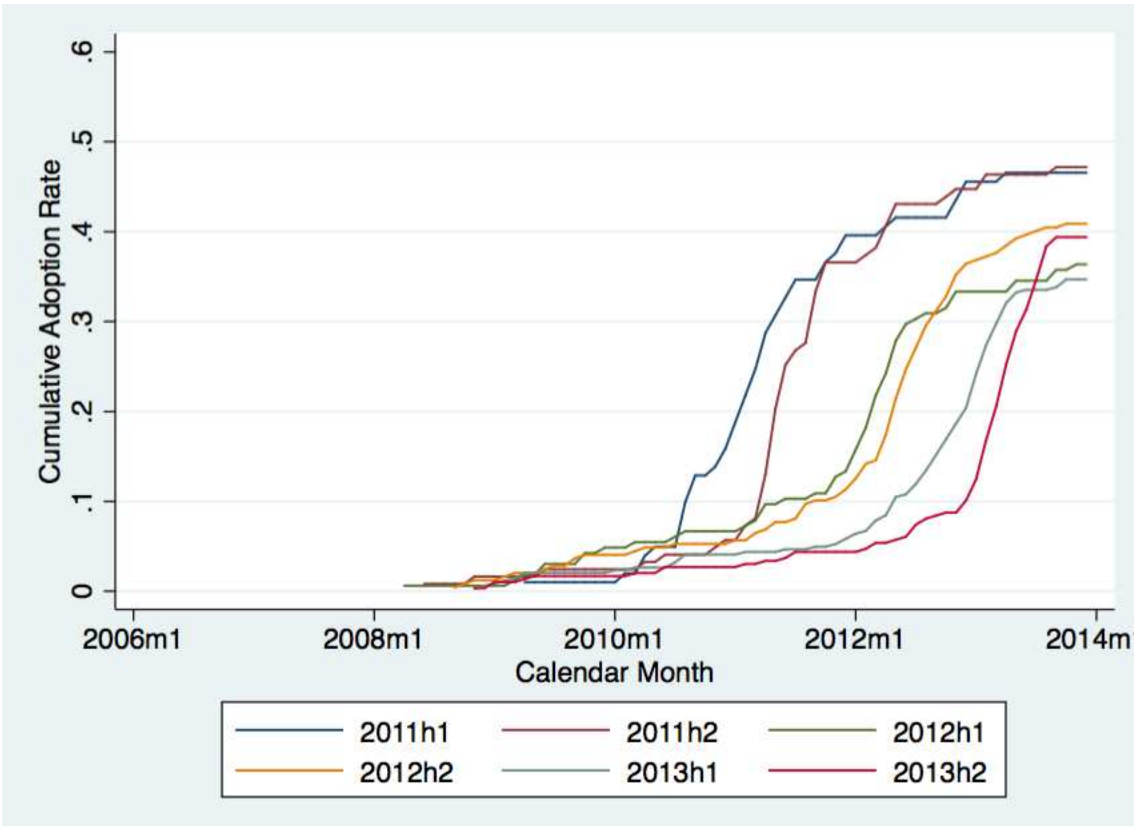

Note: This figure replicates Figure 3, but separates each generation cohort semi-annually from 2008 to 2013. For clarity, the half-year cohorts are split between Figures A3 and A4. Each line depicts s_{it} for a given cohort of shows that premiered that half-year, which is calculated as the cumulative number of shows that adopted Twitter at time t divided by the total shows in each respective cohort during the entire sample period. Each cohort includes shows that aired during that half-year.