

The sales effect of word of mouth: a model for creative goods and estimates for novels

Jonathan Beck

Received: 11 October 2006 / Accepted: 20 December 2006 /
Published online: 19 January 2007
© Springer Science+Business Media B.V. 2007

Abstract Weekly sales of creative goods—like music records, movies, or books—usually peak shortly after release and then decline quickly. In many cases, however, they follow a hump-shaped pattern where sales increase for some time. A popular explanation for this phenomenon is word of mouth among a population of heterogeneous buyers, but previous studies typically assume buyer homogeneity or neglect word of mouth altogether. In this paper, I study a model of new-product diffusion with heterogeneous buyers that allows for a quantification of the sales effect of word of mouth. The model includes Christmas sales as a special case. All parameters have an intuitive interpretation. Simulation results suggest that the parameters are estimable for data that are not too volatile and that cover a sufficiently large part of a title's life cycle. I estimate the model for four exemplary novels using scanner data on weekly sales.

Keywords New-product diffusion · Word of mouth · Creative industries

JEL codes C22 · L82 · M3

I benefitted from detailed comments by Michal Grajek, Holger Stichnoth, Christophe Van den Bulte, and an anonymous reviewer, discussions with Lars-Hendrik Röller and participants at the Vienna ACEI Conference, the Pittsburgh INFORMS Marketing Science conference and the WZB seminar, as well as proof-reading by Christopher Xitco. All remaining errors are mine. I am grateful to *Media Control GfK International* and Uwe Maisch for data provision and to the *Börsenverein des Deutschen Buchhandels* for permission to use it. I gratefully acknowledge the financial support from the German Federal Ministry of Education and Research, project *InterVal* (01AK702A), and the hospitality of the Economics group at London Business School.

J. Beck (✉)
Humboldt University and Wissenschaftszentrum Berlin (WZB), Reichpietschufer 50, 10785
Berlin, Germany
e-mail: beck@wz-berlin.de

1 Word of mouth and hump-shaped sales

Weekly sales of creative goods—like music records, movies or books—usually peak shortly after release and then decline quickly. In many cases, however, they follow a hump-shaped pattern where sales increase for some time. For example, Fig. 1 plots weekly sales for two novels released on the German book market in 2003.¹ The horizontal axis in each panel of Fig. 1 denotes calendar weeks. Sales of title 1 increase for about 7 weeks after release. Sales of title 2, instead, decrease quite steadily from its second week onwards—their sixth-week sales peak seems to be an outlier within a generally negative trend.

An additional characteristic of creative industries is that Christmas purchases often represent an important share of a title's overall sales. Figure 1 therefore highlights sales in December weeks. The Christmas effect seems to be relatively strong for title 2 and modest for title 1. Similar sales patterns can be observed for music records and theatrical movie releases (Moe & Fader, 2001; Moul, 2006), and a popular explanation for hump-shaped sales is product-specific word of mouth among heterogeneous buyers: an impressive novel can lead some readers to be enthusiastic promoters *vis-à-vis* potential buyers, such that sales increase for some time after the title's release.²

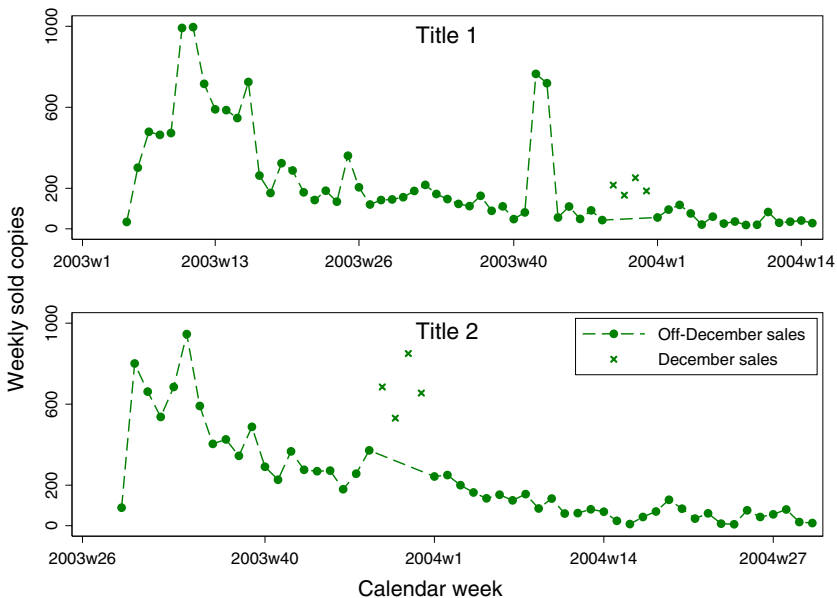


Fig. 1 Weekly unit sales for example titles

¹ Weekly sales (original hardcover edition), based on scanner data from a large sample of German retailers. Titles are anonymized; see Sect. 4 for a more detailed description of the data.

² A related economics literature explains a hump-shaped pattern in the adoption of new technologies, which is equivalent to an S-shaped pattern in their penetration, generally by either decreasing costs or increasing benefits of adoption over time (Hall & Kahn, 2003). These factors—in my application the retail price, for example—are typically constant during the short selling period of a new creative good; however, one may interpret word of mouth as evidence of changes in the (perceived) benefits of adoption.

The (non-)occurrence of word of mouth can thus determine success or failure of a particular title:

“Although *nobody knows* its fate when a new creative good appears, social contacts transmit consumers’ appraisals at a very low perceived cost to them, giving ‘word of mouth’ its importance for a creative good’s ultimate success. [...] The distribution of consumers between ‘buffs’ and ‘casuals’ strongly influences the organization of an art realm.” (Caves, 2000, p. 173)

In this paper, I study a model of new-product diffusion that features two distinct types of buyers: independent buyers and buyers influenced by word of mouth. In this respect, the model thus closely follows Caves’ description of creative industries. As a result of this form of buyer heterogeneity, the model has the realistic property that the lack of word of mouth leads to unambiguously lower sales.

Most previous studies of the sales performance of creative goods, in contrast, either assume buyer homogeneity or concentrate on other aspects than word of mouth. In addition, most of this literature has focused on motion picture success (Elberse, Eliashberg, & Leenders, 2006, provide a recent review). Regarding the effects of word of mouth in particular, De Vany and Walls (1996) present a model of information cascades with the property that box office failures and successes are equally likely and largely determined by opening week performance. The empirical section of their paper, however, focuses on characterizing overall movie demand rather than estimating title-specific parameters of word of mouth. Moul (2006) quantifies word of mouth effects in weekly box office data based on a static discrete choice model, but he does not report title-specific estimates either.

The empirical literature on the music or book industry is less extensive. Moe and Fader (2001) analyze the intertemporal pattern of music record sales as a mixture of sales to different buyer segments, but their model does not allow buyers to engage in word of mouth. Van den Bulte and Joshi (2006) use the same data to compare estimation results for a class of diffusion models with buyer heterogeneity. I discuss these models in some more detail in Sect. 2.1. A number of recent studies look at online book sales (Chevalier & Goolsbee, 2003, for example) but with focus on internet-related questions rather than book market particularities.

More relevant to this paper is the study by Chevalier and Mayzlin (2006) on a particular case of word of mouth about books: online consumer reviews. As expected, they find that positive consumer reviews tend to increase sales. Sorensen and Rasmussen (2004) use scanner data on weekly book sales—similar to the data employed in this paper—to study the sales effects of *New York Times* book reviews. They find that both positive and negative reviews tend to increase sales. Hence, both studies corroborate the view that word of mouth is important for book sales and that it comes in various forms.

This paper is complementary to these studies in that it provides estimates for the overall effect of word of mouth on a title’s sales rather than two special cases. In addition, these studies follow a reduced-form estimation approach, whereas my estimates are based on an explicit model of new-product diffusion, and thus have a structural interpretation. I also include Christmas sales as a special case of word of mouth.

Like most models of new-product diffusion, it is highly nonlinear in its parameters. In Sect. 3 of the paper, I therefore assess its estimation properties by a set of

Monte Carlo simulations. Results suggest that the parameters are estimable if the data cover a sufficiently large part of a title's life cycle and are not too volatile. In particular, book sales data often exhibit solitary outliers such as those for title 1 in weeks 42 and 43 of 2003. These outliers may be the combined effect of a demand shock, for example a television broadcast that featured the title or its author, and residual autocorrelation.

Finally, in Sect. 4, I present and discuss estimates for four exemplary novels, including those presented in Fig. 1. Estimates appear informative, intuitively reasonable and in line with 'stereotype' attributes commonly ascribed to creative goods.

The structural empirical approach proposed in this paper can be useful to applied research in various creative industry domains. For example, publishers can compare title-specific estimates to evaluate marketing strategies; applied to a single title, the model provides an estimate of consumer segmentation. A more policy-oriented analysis across titles and market segments may help characterize *ex ante* demand uncertainty in the spirit of De Vany and Walls (1996, 2004).

2 New-product diffusion with heterogeneous buyers

The idea that new-product diffusion is driven by word of mouth among heterogeneous buyers is not new. For example, a popular theory holds that the product life-cycle is determined by different adoption times of different consumer groups influencing each other, e.g., by innovators who adopt because of 'intrinsic' motivations and by imitators who adopt because the innovators do so (Rogers, 1995). Yet, most quantitative studies of new-product diffusion follow Bass (1969) and assume a homogeneous population of buyers, whose purchases are driven by 'intrinsic' and 'social' motivations at the same time.³

In principle, one could apply the Bass (1969) model to book sales data in order to generate a reduced-form estimate for the word of mouth effect; that approach, however, would miss an important feature of word of mouth in creative industries. Within a homogeneous population, word of mouth can only affect the speed of diffusion, that is, the time until every potential buyer has purchased the product. Word of mouth in Bass-type models cannot generate product failures or successes, because it has no effect on the total number of buyers. Moreover, scholars have argued that ignoring buyer heterogeneity can lead to biased estimates in applied research (Van den Bulte & Lilien, 1997; Bemmaor & Lee, 2002).

In what follows, I propose a parsimonious model of new-product diffusion with two buyer segments and argue that it represents a good tradeoff between validity and empirical tractability—especially for the analysis of aggregate sales data of new creative goods.⁴ The model coincides formally with a special case of a class of new-product diffusion models independently developed by Van den Bulte and Joshi (2006). A relevant difference is my explicit treatment of Christmas sales. I also

³ The Bass (1969) model was independently studied also in sociology, economics, as well as the statistics literature; for reviews see Mahajan, Muller, and Wind (2000), Geroski (2000) and Van den Bulte and Joshi (2006).

⁴ Applications with more detailed retailer- or consumer-level data are outside the scope of this paper; for such data, discrete choice models are more appropriate. Yet, decision makers other than retailers typically have access only to some kind of aggregated sales data (if any).

propose a more intuitive interpretation of model parameters. For convenience, I present the model mostly in terms of my application: books.

Model. Consider a new title that has just been released. Some buyers—buffs, fans of the author—are independently aware and intrinsically motivated to buy this title. In addition, they may act as promoters *vis-à-vis* other buyers, and are therefore labeled ‘influentials.’⁵ A second type of buyer—the ‘imitator’—only buys after being informed by word of mouth from at least one influential. Such word of mouth can take various forms, for example personal recommendations or book reviews. In fact, it may not involve a lot of talking: merely observing an influential person reading a certain book may induce imitator interest.

Denote with M the total number of potential buyers of the title under consideration. M does not include multiple or repeat purchases, total sales may therefore exceed M ; below, I introduce the particular case of Christmas present purchases by previous buyers. Let θ denote the share of influentials in M . Sales to both influentials and imitators are distributed over time, but differently. Denote by $F_i(t)$ the c.d.f. of sales to group i ($i = 1,2$)—that is, cumulative sales to group i divided by its population—and by $f_i(t)$ the corresponding density.

Not all influentials buy the title straightaway in its release week, for example for budgetary reasons or lack of immediate leisure time. A standard assumption, instead, is that in every period the title is bought by a fraction p of those independent buyers who have not bought earlier. Excluding potential word of mouth effects, such behavior generates the steady decay pattern common to box office figures. In continuous terms, p is a hazard rate: $p = \frac{f_1(t)}{1-F_1(t)}$. Since $F_1(0) = 0$, the cumulative number of sales to influentials at time t , can be solved for as

$$n_1(t) = \theta M F_1(t) = \theta M (1 - e^{-pt}). \tag{1}$$

The number of imitator purchases depends on whether there is positive word of mouth by influential buyers. In particular, every influential recommends the product to \tilde{q} imitators each period following her purchase ($\tilde{q} \geq 0$). If the title is not found worth recommending ($\tilde{q} = 0$), sales are determined by Eq. (1) and thus limited to a share θ of their potential M . Notice that the model does not allow for negative (sales-destructive) word of mouth; I return to this point in the next section.

If $\tilde{q} > 0$, the contacted imitators go ahead and buy the recommended title, unless they have not already done so due to an earlier recommendation.⁶ Imitators do not recommend the product to other people (but they may offer it as a Christmas present, see below). Since the probability that an imitator exposed to word of mouth at time t has not been contacted and therefore has not bought earlier is $1-F_2(t)$, the cumulative number of sales-effective recommendations at time t is $(1-F_2(t))$ times \tilde{q} times $n_1(t)$. Divided by the overall number of imitators $(1-\theta)M$ this amounts to the density $f_2(t)$ and rearranging yields the relationship

⁵ The literature on new-product diffusion often focuses on technical innovations such as color television or the washing machine and therefore labels this consumer segment ‘innovators.’ Here, I follow Keller and Berry (2003) and Van den Bulte and Joshi (2006) in using a more generally applicable term.

⁶ A formal assumption is that, at each point in time, each influential contacts a different set of imitators. In a continuous-time model, the probability that an imitator is contacted by more than one influential at once should be negligible.

$$\frac{f_2(t)}{1 - F_2(t)} = qF_1(t), \quad (2)$$

where $q = \tilde{q} \frac{\theta}{1-\theta}$ is a convenient reparametrization, and $F_1(t) = \frac{n_1(t)}{\theta M}$. Using Eq. (1) and the fact that $F_2(0) = 0$, this differential equation solves for

$$F_2(t) = 1 - e^{\frac{q}{\theta}(1-e^{-pt}-pt)}. \quad (3)$$

Hence, cumulative sales to imitators are $n_2(t) = (1-\theta) M F_2(t)$, and we get for total cumulative sales at time t :

$$N(t) = n_1(t) + n_2(t) = M(1 - \theta e^{-pt} - (1 - \theta)e^{\frac{q}{\theta}(1-e^{-pt}-pt)}). \quad (4)$$

Interpretation. It is possible that a particular title does not stimulate word of mouth ($q = 0$), or that the group of imitators has zero mass ($\theta = 1$). In either case, the model reduces to a negative exponential growth model and total sales $N(t)$ are equal to $n_1(t)$ as defined by Eq. (1).⁷ Other boundary values for θ and p are theoretically impossible or highly improbable for real data. A title with $p = 0$ or $\theta = 0$ would never exhibit positive sales, and a title with $p = 1$ would sell to influentials only in week 1.

Obviously, the proposed 2-segment structure is a modest form of buyer heterogeneity: after all, everybody's different. Yet, the diffusion parameters also have an interpretation as sub-population averages that allows for some unsystematic within-segment heterogeneity. For example, an implicit assumption is that influential buyers immediately evaluate a title and decide whether to recommend it. That assumption is not overly restrictive, because readers often evaluate a novel already from its first couple of pages. Nonetheless, by introspection I presume that some individuals are inclined to irrationally pile purchased books on their bedside table for years before reading and eventually recommending them. In this case, q can be regarded as an average recommendation rate: some influentials recommend more, some less, and the weekly average is q .

Identification. Equation (4) illustrates that the only 'variable' in this and similar models of new-product diffusion is time (t). Correspondingly, the parameters of the model are identified by the shape of the sales curve: if there is positive word of mouth ($q > 0$), sales follow a hump-shaped pattern over time and the particular shape of this curve, combined with the assumptions of the model, allows recovery of the respective parameter values.

Yet, Eq. (4) also illustrates that M and θ are not identified if $q = 0$ and that q is not identified if $\theta = 1$; that is, if sales are not hump-shaped but only follow the decay pattern of Eq. (1). The latter case—positive word of mouth without recipients—is arguably unrealistic. But the former case lies at the heart of marketing problems in creative industries: you never know how much a title would have sold had it had some word of mouth!

Therefore, if estimated parameters include $\hat{\theta} \approx 1$ and/or $\hat{q} \approx 0$, the other estimates are not reliable and a more restricted model may be applied to the respective

⁷ Notice that although the Bass model also reduces to Eq. (1) if its word of mouth parameter equals zero, the two models are not nested.

data. Estimates for p or θ that are unreasonably close to zero might suggest misspecification (systematic unobserved effects), or simply that sample size is too small or error variance too large. In Sect. 3, I assess the importance of the latter two effects with a Monte Carlo simulation.

2.1 Possible extensions

The proposed model represents one out of several possibilities to formalize new-product diffusion within a heterogeneous (2-segment) population. At least for an application to creative goods markets, it seems reasonable to assume that (i) only influential buyers engage in word of mouth, which (ii) is directed only towards imitator buyers. Note that these assumptions only refer to sales-effective communication; they do not exclude general communication about a certain title that has no direct sales effect. Most readers like to talk about a book they have read, and they use forums like *Oprah's Book Club* to coordinate reading and follow-up communication. Although this holds for influential buyers (buffs) as well—and may affect the timing of their purchase (parameter p)—it does not affect the purchase decision. Imitators (casuals), instead, buy after and as a consequence of communication. In their case, word of mouth is sales-effective, and it is this particular word of mouth I focus on.

Negative word of mouth. One restriction in the above model is that it only considers positive word of mouth from influentials to imitators, although anecdotal evidence from the movie industry suggests that negative (sales-destructive) word of mouth can exist as well. In the 2-segment framework, negative word of mouth implies not only that $q = 0$, but in addition that specific communication among influentials dissuades some of them from their initially planned purchase. As Eq. (1) illustrates, this would introduce another kind of identification problem: with negative word of mouth among influentials, their number (θM) would be subject to change over time.

Based on aggregate sales data alone, however, it is impossible to distinguish this effect from the decay parameter p . As an example, consider title 2 in Fig. 1, whose sales start declining soon after release. Based only on the figure, it is impossible to say whether the quick decline is due to negative word of mouth (decreasing θM over time) or a large decay parameter p . In order to identify an effect of negative word of mouth, one would either need more detailed data on post-purchase communication, such as online consumer reviews, or be willing to impose restrictions on p in Eq. (1). Moul (2006) basically follows the latter approach, yet on the basis of a static discrete choice model.

Word of mouth among imitators. In the above model, imitators care exclusively for recommendations by influentials. In general, however, word of mouth among imitators may also play a role. Van den Bulte and Joshi (2006) analyze a more general class of models which allows for positive word of mouth within both segments. These more general cases have the property that period sales can exhibit a 'dip' in early sales periods, when decreasing sales to influentials are not yet over-compensated by increasing sales due to word of mouth.

The data studied in this paper do not seem to exhibit such patterns. For similar data from the U.S. market, Sorensen (2006) finds that sales for most titles peak in early sales weeks (before the fifth week), which is inconsistent with the 'dip' property

of the more general models analyzed by Van den Bulte and Joshi. Whereas Sorensen (2006) concludes that his econometric assumption of a monotone sales pattern over time is not critical for his purposes, the approach followed here is to explicitly consider word of mouth as the factor underlying a potential non-monotone pattern in early weeks.

Although attractive from a theoretical viewpoint, the more general models by Van den Bulte and Joshi (2006) come at the cost of at least one additional parameter. From an empirical viewpoint this is a drawback, as estimation difficulties usually increase exponentially with the number of parameters to be estimated. Moreover, the respective cumulative sales function—the analog to Eq. (4)—becomes rather complex and estimation based on standard methods appears to be troublesome. Reported estimates based on the most general model and data for 19 music CDs are often insignificant or seem to converge to theoretical limit values (0 or 1).

Whereas unsatisfactory statistical significance may be an inevitable consequence of large data variance, degenerate estimates suggest that simpler models may suffice to analyze new-product diffusion in creative industries. Put differently, the nature of the data may often not allow the identification of flexible diffusion models, in which case researchers have to make structural assumptions. For the purposes of this paper, the 4-parameter model presented above—which already provides for sufficient estimation issues (cf. Sect. 3)—appears as a reasonable approximation to word of mouth communication about novels. Nevertheless, a detailed empirical comparison between alternative models, beyond that provided by Van den Bulte and Joshi (2006), is a promising subject of further research.

2.2 Incorporation of Christmas sales and estimation

So far I have neglected Christmas sales, which can represent an important share of overall sales. I propose the following Christmas effect. Within a certain time period—here I let that be the month of December—every previous purchaser of the title buys c extra copies in order to give them away as presents. Alternatively, one may assume that only influentials or only imitators purchase extra copies as Christmas presents; estimation and interpretation of parameter c would be modified in a straightforward manner (see below).

Christmas presents can be interpreted as a special case of word of mouth: instead of expressing a recommendation to read in a conversation, the recommender goes one step forward and assumes her counterpart's purchase decision.⁸ Whereas such recommendation-by-donation directly leads to increased sales of a title, it does not necessarily lead to increased reading. The donor may be mistaken in the donee's overall interest to read the presented title or in her interest to read it at that moment.⁹ I therefore assume that copies purchased as Christmas presents do not instantaneously lead to additional word of mouth and corresponding sales. This is also consistent with the model's assumption discussed earlier, namely that there are no second-order word of mouth effects through communication among imitators.

⁸ Exemplary cases are reported in 'Women buy fiction in bulk and publishers take notice', *New York Times*, March 17, 1997.

⁹ The potential "deadweight loss" involved with misguided Christmas presents is subject to an academic debate initiated by Waldfogel (1993).

Observed December sales then consist of two unobserved components: (i) the usual sales depending on the diffusion parameters and (ii) additional Christmas sales. A potential extension would be to allow Christmas presents to affect word of mouth in sales weeks *after* Christmas, when the donees have had a chance to unwrap their present and actually read the book. In that case, sales would pick up again after Christmas, relative to November, for example. At least for my example titles, however, this does not seem to be the case.

Furthermore, there is no reason to expect a particular distribution of the additional Christmas sales across the December sales weeks. For example, sales of title 1 in Fig. 1 seem to exceed their general trend in all four December weeks and by a similarly modest amount. Based on this observation, one might model Christmas sales such that December sales enjoy a linear markup on top of the ‘usual’ sales predicted by the model parameters for that period. December sales of title 2, however, vary strongly across the four weeks. Therefore, I do not impose a particular pattern but rather assume that Christmas sales are randomly distributed across the four December weeks. Sources for such randomness could be title-specific differences in store availability or aptness to be last-minute presents.

Estimation. The standard approach to estimating the parameters of a new-product diffusion model is based on period sales, that is, first differences of the cumulative sales function (Putsis & Srinivasan, 2000):

$$S(t) = N(t; \phi) - N(t - 1; \phi) + \sum_{i=1}^4 \lambda_i D_{ti} + \varepsilon_t, \tag{5}$$

where $S(t)$ denote sales of a given title during the period $(t-1, t)$, $N(\cdot)$ is defined by Eq. (4), ϕ is the set of model parameters (M, θ, p and q), ε is a stochastic error term, and $t = 1, \dots, T$.¹⁰

The D_{ti} are dummy variables that account for the random distribution of Christmas sales across December weeks. More precisely, D_{t1} is equal to one if period t is the first week of December, and D_{t2}, D_{t3} , and D_{t4} are defined analogously. Constrained to be non-negative, these dummies’ coefficients (λ_i) capture all December sales that exceed those predicted by the otherwise best-fitting set of parameter values.¹¹

Parameter c is implicit in the estimated λ -coefficients. Its exact estimate depends on whether one wants to impose a certain minimum delay between first purchase and Christmas gift decision. In other words, one may want to assume that only buyers who bought a title for own reading before week t^C qualify as potential Christmas donors. The additional Christmas sales are then c times $N(t^C; \phi)$, and an estimate for c is

¹⁰ In some applications—including mine, as I discuss in more detail in Sect. 4—it may be difficult to determine the exact release date for the title under consideration. First- and second-week sales figures are then hardly comparable and estimation should be adapted to this problem.

¹¹ Parameter identification would be complicated if December observations coincided with the early, increasing part of a title’s sales curve. In practice, the vast majority of popular books is released with a safe distance to Christmas between spring and early fall. Presumably, the fact that the book season starts with the book fair in October is an effort to animate word of mouth early enough such that it can boost Christmas sales.

$$\hat{c} = \frac{\sum \hat{\lambda}_i}{N(t^C, \hat{\phi})}. \quad (6)$$

In my application below, t^C is the last week of November; I thus assume that December for-read purchases are too short-term to lead to additional for-present purchases.

Residual autocorrelation. It is possible that a shock in one week's sales figure has repercussions in following weeks. An example would be a television broadcast featuring the title or its author that leads to increased sales not only in the week of the event but also in the following week, because some viewers are slower in executing their shopping list. The error term ε would then follow an AR(1) process; the twin outliers in weeks 7/8 and 42/43 for title 1 seem to indicate a case in point. Such residual autocorrelation is conceptionally different from—and may arise on top of—the serial correlation of model sales inherent to the model diffusion process as captured by t in the estimation equation. The econometric procedure to test and account for autocorrelation in diffusion models is quite standard and follows Frances (2002); see Beck (2006b) for some more detail.

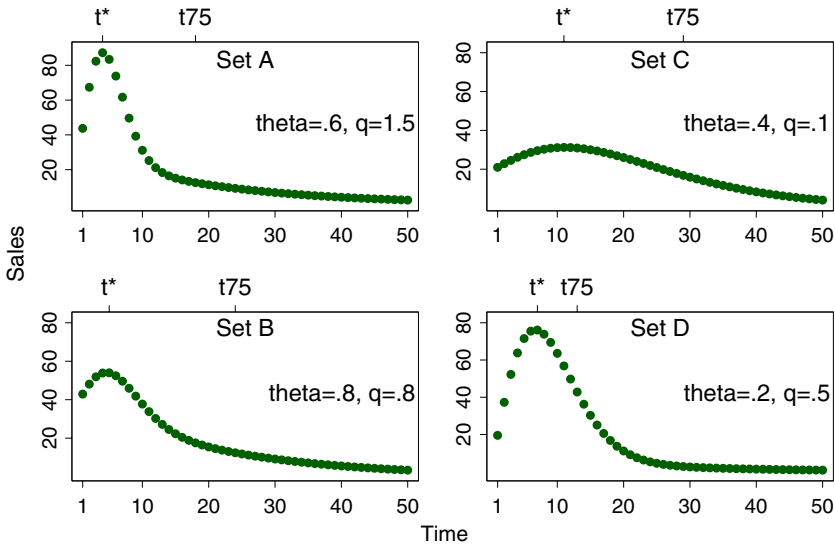
The parameters of Eq. (5) can be estimated by nonlinear least squares (NLS). A grid search procedure, evaluating the sum of squared residuals for different combinations of parameter values, is used in order to find proper initial values for iterative estimation. Through log-transformations, I impose non-negativity for all parameters, and $\theta, p \leq 1$. Unfortunately, neither asymptotic nor small-sample properties of such NLS estimators are known (Boswijk & Franses, 2005); but their bias and consistency may be studied by means of a Monte Carlo simulation. See Van den Bulte and Lilien (1997) and Bemmaor and Lee (2002) for simulation studies of the Bass (1969) model.¹²

3 Simulation

For an assessment of the properties of the proposed estimation strategy, I use Eq. (4) to generate artificial data for four different sets of parameters and $T = 60$. Only the parameters θ and q vary across sets. In all four sets, $p = .05$, $M = 1,000$ and $c = 0$. Since the Christmas effect is additive, its exclusion should not affect the validity of the simulations. I chose parameter sets (θ, q) such that they represent different stereotypes of titles: the 'shockseller' (set A), the 'sleeper' (set C), and two intermediate types (B, D). Parameter sets A, B, and C are also similar to estimates based on real data as reported in Sect. 4. Figure 2 summarizes parameter values and presents the resulting sales patterns for each set. The peak sales date of the respective series is indicated by t^* ; t_{75}^* indicates the point in time where the series reaches 75% of total sales.

Following Van den Bulte and Lilien (1997), I perturb the artificial data with a multiplicative, log-normally distributed AR(0) error term u_t that has an expected

¹² Boswijk and Franses (2005) propose an alternative estimation method for the Bass model based on weighted least squares and argue that it has better properties than the NLS approach, which has been implemented widely in similar applications. However, their approach is not applicable to the proposed diffusion model in a straightforward manner.



In all sets, $M=1000$ and $p=.05$

Fig. 2 Simulated sales patterns

value of 1 and a variance of $\exp(\sigma^2)-1$.¹³ I assess three different error variances, where σ equals .06, .42 or .78, corresponding to the lowest, the intermediate and the highest variance used by Van den Bulte and Lilien (1997). For example, $\sigma = .06$ implies that with a probability of about 95%, the difference between artificial and perturbed data is not larger than $\pm 10\%$. For $\sigma = .78$, instead, this difference is likely to lie within -80% and $+200\%$. Comparing simulation results for different error variances demonstrates how sensitive estimates are to solitary shocks of different intensity.

I estimate Eq. (5) with the perturbed data for $S(t)$, using the first 24, 40 or 56 observations. Estimates from diffusion models are usually quite sensitive to sample size; in particular, smaller samples that do not include t^* make estimation of M unfeasible (Van den Bulte and Lilien, 1997). For each simulation scheme, consisting of a parameter set (4 variants), error variance (3 variants) and sample size (3 variants), the perturbation and estimation procedure is repeated 1,000 times. Altogether, this yields 36,000 separate estimation results.

As error variance increases, a number of simulations result in degenerate estimates at or very close to boundary values zero (p and θ) or one (θ). In these cases, probably large absolute errors occurred in the early observations of the respective sample. Imagine large shocks in the first ten sales weeks of a type-B title as depicted in Fig. 2. Any estimation procedure may then have difficulties in identifying a hump-shaped sales curve and may converge to a set of parameter values that imply negatively sloped sales, like $\theta = 1$ or $p = 1$. Effectively, degenerate estimates indicate that the data are such that the procedure cannot distinguish this model from a simpler, 2- or 3-parameter model.

¹³ If v_i is normally distributed with mean $-\frac{\sigma^2}{2}$ and variance σ^2 , then $u_i = \exp(v_i)$ is log-normally distributed with mean 1 and variance $\exp(\sigma^2)-1$.

Table 1 Share of degenerate results across simulation sets

	Set A	Set B	Set C	Set D
<i>Sample size: T = 24</i>				
$\sigma = .06$	0	0.09	0.002	0
$\sigma = .42$	0.25	0.16	0.52	0.62
$\sigma = .78$	0.54	0.50	0.76	0.72
<i>Sample size: T = 40</i>				
$\sigma = .06$	0	0.001	0	0
$\sigma = .42$	0.07	0.06	0.13	0.45
$\sigma = .78$	0.33	0.30	0.44	0.62
<i>Sample size: T = 56</i>				
$\sigma = .06$	0	0	0	0
$\sigma = .42$	0.02	0.04	0.06	0.41
$\sigma = .78$	0.21	0.19	0.32	0.57

Therefore, the degree to which the occurrence of degenerate results varies across simulation schemes is an indication of the estimator's sensitivity to differences in sample size and data variance. For a comparison across schemes, I define the estimation results of any repetition as degenerate if they fulfill at least one of the following conditions: (i) \hat{p} is smaller than .01 and not significantly different from 0 (95% confidence), (ii) $\hat{\theta}$ is smaller than .025 or larger than .975 and not significantly different from 0 or 1, respectively, (iii) \hat{M} is more than tenfold observed cumulative sales.¹⁴ Correspondingly, Table 1 lists the share of degenerate results within the 1,000 repetitions for each of the 36 simulation schemes. Estimation difficulties seem to be mainly driven by data variance (σ): degenerate estimates are rare as long as $\sigma = .06$. With increasing σ , all parameter sets are prone to yield degenerate results, but the effect is weaker for large samples.

A positive side to the occurrence of degenerate p - or θ -estimates is that they are easily identified in practice: they are coupled with slow convergence and unreasonably large estimates for q (in the thousands) and/or M (in the millions). Given such results, the analyst may attempt to reduce data volatility before re-estimation, for example by identifying the specific causes for solitary outliers.

Non-degenerate simulation results, on the other hand, seem fairly accurate. Table 2 compares means and medians of the parameter estimates for two simulation schemes: one with low error variance ($\sigma = .06$) but based on a relatively small sample ($T = 24$), and one with both intermediate error variance and sample size ($T = 40$, $\sigma = .42$). Especially with low error variance and for sets A, B, and D, estimates seem to be close to their true values both in terms of the average and the median.

Estimation seems to be more difficult for set C, however, there the first 24 observations cover a much smaller share of overall sales than for the other sets. One conclusion from this finding is that estimation of the model requires data on the major share of a title's first-release life cycle, and therefore, that the model's forecasting capabilities are limited; I return to this point below.

For every parameter set, increased error variance seems to induce a non-negligible bias on one or the other parameter estimate (mostly q). Yet, this finding is in line with previous simulation studies, and hence not specific to the model studied

¹⁴ This particular cutoff value for \hat{M} is not pivotal, since most degenerate sets of estimates come with M -estimates in the 6-digit range or higher.

Table 2 Means and medians of simulated estimates for ($T = 24, \sigma = .06$) and ($T = 40, \sigma = .42$)^a

Set	Simulation scheme		Parameter			
			M	θ	p	q
A	$T = 24, \sigma = .06$	<i>True value</i>	1,000	0.6	0.05	1.5
		Mean estimate	1,003.3	0.602	0.050	1.521
		Median estimate	999.7	0.601	0.050	1.509
	$T = 40, \sigma = .42$	Mean estimate	1,041.4	0.631	0.050	2.066
		Median estimate	1,012.2	0.617	0.047	1.769
B	$T = 24, \sigma = .06$	<i>True value</i>	1,000	0.8	0.05	0.8
		Mean estimate	1,001.0	0.800	0.050	0.806
		Median estimate	999.9	0.799	0.050	0.800
	$T = 40, \sigma = .42$	Mean estimate	1,037.7	0.738	0.060	1.155
		Median estimate	1,020.4	0.759	0.047	0.868
C	$T = 24, \sigma = .06$	<i>True value</i>	1,000	0.4	0.05	0.1
		Mean estimate	1,173.9	0.432	0.058	0.154
		Median estimate	1,113.7	0.390	0.051	0.097
	$T = 40, \sigma = .42$	Mean estimate	1,087.1	0.393	0.073	12.470
		Median estimate	1,064.1	0.456	0.044	0.100
D	$T = 24, \sigma = .06$	<i>True value</i>	1,000	0.2	0.05	0.5
		Mean estimate	1,021.8	0.201	0.055	0.559
		Median estimate	1,003.0	0.208	0.050	0.496
	$T = 40, \sigma = .42$	Mean estimate	1,008.1	0.233	0.062	0.630
		Median estimate	996.0	0.242	0.051	0.492

^a Based on 1,000 repetitions for every set of parameters. Statistics exclude degenerate results

here. Van den Bulte and Lilien (1997) as well as Bemmaor and Lee (2002) obtain similar results for the less complex 3-parameter Bass diffusion model. In particular, the results for set C suggest that a relatively flat hump-shaped trend becomes difficult to identify with increasing data variance.¹⁵

In summary, this simulation exercise shows that the model is estimable if the data include a sufficiently large number of not too volatile observations. When using volatile data, analysts should be prepared to obtain degenerate estimates and should otherwise allow for slight bias when interpreting non-degenerate estimates. Moreover, there is certainly room for methodological improvement. For example, the simulated error is a multiplicative one, which introduces heteroskedasticity depending on the absolute level of sales, whereas estimation assumes additive errors of constant variance. Yet, level-dependent heteroskedasticity is probably a realistic assumption that may be exploited to improve estimates in the fashion recently demonstrated for the Bass model by Boswijk and Franses (2005). Progress along these lines, however, is outside the scope of this paper.

4 Application

On the basis of the available number of sales observations and data volatility, I selected four titles from a sample of novels released in Germany in 2003 as

¹⁵ More detailed simulation results, such as kernel density plots of the distribution of estimates, can be found in the appendix of Beck (2006b).

hardcover editions.¹⁶ The data has been kindly provided by *media control GfK International*, a marketing research firm that collects scanner data from over 750 points of sale (booksellers, department stores) as well as all main internet retailers in Germany. Not sampled are direct sales from publishers to consumers, book club sales and mail order sales. Supermarket sales are also not sampled, but they represent a small portion of book sales.

The sampled retail channels account for about 66% of total book sales in Germany, however, for the particular segment studied here (novels in hardcover), sales coverage of sampled channels is likely to be much higher. First, publisher direct sales are not very important for popular publications such as novels; direct sales usually concern professional publications. Second, book clubs can be regarded as a secondary market that only becomes important for a title after its diffusion in the primary market (which is studied here).

Furthermore, no particular estimation bias arises from this type of sampling. By law, book prices are the same across all retail channels.¹⁷ In theory, independent buyers may be more inclined to order directly from publishers because they do not need retailer advice; in practice, however, title availability is high and ordering processes are quicker (typically overnight) at brick-and-mortar booksellers and online retailers. Direct orders from publishers are thus unattractive for non-professional buyers. Based on a specific population weight for every point of sale, the scanner data were aggregated to represent nationwide sales. The sample period ends in summer 2004, providing up to 81 weekly observations per title. As demanded by the data proprietor, I received anonymized data where all title-, author- and publisher-specific information except for sales (by week) and price (constant) had been removed.¹⁸

In contrast to theatrical movies, new books do not have a particular weekday for release. A title may be shipped on a Friday, leaving only one or two sales days in its first calendar week. The analyst may then wrongfully interpret increasing sales in the title's second calendar week as a result of word of mouth, while in fact first- and second-week figures are not comparable due to a different number of sales days. Low sales figures in the first week, as compared to later weeks, for titles 1, 2, and 3 may indicate that this is an issue for the sample titles studied here. Since I do not have information on title-specific release dates, my estimations therefore only use observations from the second week onwards, hence $t = 2, \dots, T$.

Direct estimates. As expected, after an initial estimation round based on Eq. (5) and following the procedure outlined by Franses (2002), the null hypothesis of AR(0) errors is rejected against the AR(1) alternative for title 1. For the other titles, it cannot be rejected. Therefore, I present results based on Eq. (5) for these titles. For title 1, I re-estimate the model with the AR(1) specification and for $t = 3, \dots, T$. Separately for each of the four titles, Fig. 3 plots sales and model predictions from

¹⁶ Many titles in this random sample have low overall sales and zero sales in many observed weeks, but a significant number of titles exhibits hump-shaped sales.

¹⁷ In many European countries including Germany, book prices are by law subject to resale price maintenance; that is, retailers must not offer discounts from the publisher's list price. See Canoy, van der Ploeg, and van Ours (2006) and Beck (2006a) for more details.

¹⁸ According to the raw data, single copies of title 1 and 3 were sold two and three weeks before sales of the respective title took off (with 33 and 129 copies sold in that week, respectively). I treated these early sales as erroneously booked advance orders and hence assumed that 34/130 copies were sold in the respective title's release week.

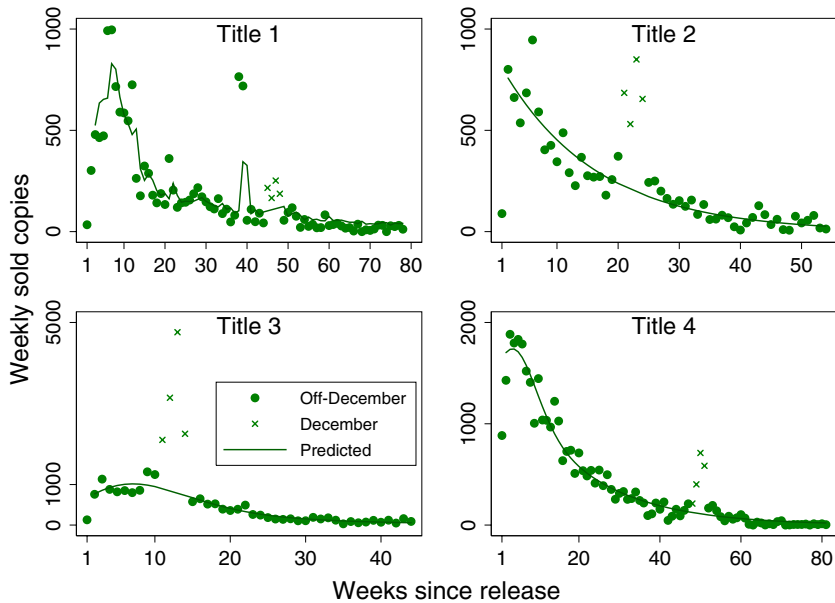


Fig. 3 Observed sales and model prediction

estimation results. Due to the incorporation of autocorrelated errors, the predicted sales pattern for title 1 is not as smooth as for the other titles.

Table 3 presents the corresponding parameter estimates. Title 1’s ρ -estimate is significant and suggests that error autocorrelation is modestly positive. The fit

Table 3 Estimation results^a

	Title 1	Title 2	Title 3	Title 4
M	17,285.027 (3823.504)	13,251.633 (692.085)	20,388.379 (1351.369)	34,350.723 (871.083)
θ	0.696 (0.094)	0.979 (.)	0.550 (0.099)	0.870 (0.099)
p	0.022 (0.015)	0.064 (0.005)	0.055 (0.014)	0.054 (0.003)
q	1.273 (0.788)	2,096.171 (.)	0.287 (0.104)	0.635 (0.114)
λ_1	137.481 (129.880)	461.750 (80.566)	1,196.847 (126.896)	88.205 (97.366)
λ_2	0 (.)	321.662 (80.549)	2,289.711 (126.160)	284.724 (97.323)
λ_3	132.403 (129.893)	653.708 (80.526)	3,965.155 (125.590)	600.899 (97.282)
λ_4	38.084 (130.816)	470.941 (80.495)	1,517.897 (125.224)	479.749 (97.241)
ρ	0.356 (0.113)	0	0	0
Observations (weeks)	76	53	43	80
R^2	0.838	0.957	0.991	0.981
Root MSE	128.3	78.9	117.9	96.5
Iterations	27	45	17	17
Total sales:				
observed	14,842	13,879	27,809	34,253
predicted	17,593	15,160	29,358	35,804
Price (EUR)	19.90	10.00	24.00	19.90

^a NLS estimates. Standard errors in parentheses (computed using the delta method, based on asymptotic standard errors for the log-transformed NLS parameters)

measures R^2 and root mean squared error (MSE), however, indicate that the data are in general more variant for title 1 as compared to the other titles, thus not only as a result of autocorrelation.¹⁹

Except for title 2, all main parameters appear reasonable and precisely estimated. With a θ -estimate very close to one, a very large q -estimate and missing standard errors, the estimates for title 2 appear to be of the degenerate kind encountered in the simulation. Not surprisingly, the present model seems to be too complex for this title, whose weekly sales decrease quite steadily from week 2 onwards. In response, the estimation procedure converges to a set of parameters which implies this steady decrease.

Regarding the estimates' economic significance, titles 3 and 4 both have an estimated period purchase probability (\hat{p}) for influential buyers of around 5.5%. This implies that it takes around 13 weeks until half of the influential buyers have purchased the title. It takes around 31 weeks with a hazard rate of 2.2% as estimated for title 1. Accordingly, by its 76th sales week, title 1 has reached a lower share of predicted total sales (sum of \hat{M} and $\hat{\lambda}_i$) than titles 3 and 4 in 43 and 80 sales weeks, respectively. Notice that, as predicted cumulative sales never reach M in finite time, the sum of \hat{M} and $\hat{\lambda}_i$ will always be somewhat higher than observed sales.

A related issue is the model's limited forecasting capacity. Clearly, from a practitioner's viewpoint it is tempting to use the model, say, after the first 26 weeks in order to obtain an estimate for the remaining sales. Yet, the simulation results imply that this will yield reasonable results only for titles with little to forecast, that is, for titles whose bulk of sales have already occurred. Accordingly, estimations based on the first 27 weeks since release yield relatively accurate predictions for titles 2 and 4, whereas for titles 1 and 3, they fail to converge or return degenerate results (details omitted). In any case, Christmas sales cannot be forecast without restrictive assumptions on c .

The estimated word of mouth parameters seem to suggest some stereotype attributes for the observed titles. Title 1, for example, resembles the shockseller pattern of simulation set A: influentials make up for a comfortable majority of buyers ($\hat{\theta} \approx .7$); combined with a large word of mouth coefficient \hat{q} , sales quickly and strongly peak, but then resume to a relatively slow decay pattern from week 20 onwards. Title 3, in contrast, with a relatively low q -estimate tends to have sleeper qualities, leading to sales that decrease much more slowly over time. Sales of title 4 exhibit an early but moderate peak. Correspondingly, its θ -estimate suggests that a relatively small share of its buyers were driven by word of mouth (13%, excluding repeat Christmas purchases).

Indirect estimates. In order to account for Christmas present purchases as a special case of word of mouth, I use the λ -estimates and information on pre-December predicted cumulative sales $N(t^C, \hat{\theta})$ —where t^C denotes the last November sales week—to obtain an estimate for c (cf. Eq. (6)). Also, recall the word of mouth coefficient \tilde{q} from Sect. 2, which has a more intuitive interpretation than q and is obtained from estimates \hat{q} and $\hat{\theta}$ (cf. Eq. (2)). Table 4 summarizes both c - and \tilde{q} -estimates, as implied by estimation results.

For title 1, the estimates \hat{q} and $\hat{\theta}$ imply a \tilde{q} of about .6, indicating that almost six out of 10 influential buyers are estimated to have recommended title 1 to imitators

¹⁹ I include the R^2 measure only to compare fit across titles. High R^2 values are common in nonlinear models and not per se suggestive of a good specification (Trajtenberg & Yitzhaki, 1989).

Table 4 Indirect estimates

	Title 1	Title 3	Title 4
\tilde{q}	0.556	0.234	0.095
c	0.024	1.02	0.047
$N(t^c, \hat{\theta})$	13,012	8,796	31,107
$\sum \hat{\lambda}_i$	308	8,969.6	1,453.6

each period. A \hat{c} of about .02 suggests that on average only 2% of all pre-December buyers of title 1 are estimated to have purchased an additional copy as Christmas present in December. The other two titles have lower estimates for \tilde{q} , but higher \hat{c} . Title 3 is estimated to have generated about one Christmas present purchase per pre-December buyer on average. Interestingly, a negative correlation between \tilde{q} - and c -estimates would be suggestive of a particular typology of book recommendations: some titles are prone to instantaneous recommendation through standard word of mouth (as measured by a high \tilde{q}), while for other titles standard word of mouth is withheld (low \tilde{q}) in favor of a stronger form of recommendation—a surprise Christmas present (high c).

5 Conclusion

Conventional wisdom holds that word of mouth between heterogeneous buyers is a crucial success factor for new creative goods. In this paper, I study a model of new-product diffusion that captures some of the particularities of word of mouth on new creative goods and especially the book market. Important features of this model are a heterogeneous population of buyers, whose word of mouth can have an effect on total sales of a title, and the explicit treatment of Christmas present purchases.

Like most diffusion models, it is highly nonlinear in its parameters, which can complicate estimation in certain circumstances. A simulation exercise shows that parameters can be estimated fairly accurately with sufficiently many, not too volatile observations. Estimation results for four sample titles appear informative and intuitively reasonable and thus illustrate the model's applicability. Development of empirical methodology is warranted, however, in order to further reduce estimation bias.

Both academic and professional researchers can fruitfully apply the approach I present in this paper. For example, an *ex-post* estimate of consumer segmentation for a particular hardcover edition of a title can inform publishers' marketing strategies regarding later editions of the same title (paperback, book club, etc.) as well as for future titles by the same author. In an analysis of sales of a single title, information on events behind outliers can be added to the model in the same fashion as the Christmas dummies in order to estimate the additional sales effect of these events. Publishers can compare estimates across titles in order to test and evaluate marketing strategies.

Researchers concerned with public policy may also benefit. Demand uncertainty is often claimed to be an essential characteristic of creative industries: the *nobody knows* property (Caves, 2000). Regarding the book industry, demand uncertainty

serves as a building block to justify the use of vertical restraints such as resale price maintenance (in Europe) or returns policies (in the U.S.). The micro-economic foundations of demand uncertainty in creative industries are, however, in general unclear. This paper highlights a clear form of demand uncertainty: uncertainty about the occurrence and intensity of positive word of mouth. Applied to a larger number of titles, its empirical approach to word of mouth could shed light on the global nature of demand uncertainty in the book trade.

References

- Bass, F. M. (1969). A new product growth model for consumer durables. *Management Science*, 15(5), 215–227.
- Beck, J. (2006a). Fixed, focal, fair? Book prices with optional resale price maintenance. Working paper.
- Beck, J. (2006b). The sales effect of word of mouth: A model for creative goods and estimates for novels. WZB Discussion Paper SP II 2006-16.
- Bemmar, A. C., & Lee, J. (2002). The impact of heterogeneity and ill-conditioning on diffusion model parameter estimates. *Marketing Science*, 21(2), 209–220.
- Boswijk, H. P., & Franses, P. H. (2005). On the econometrics of the Bass diffusion model. *Journal of Business & Economic Statistics*, 23(3), 255–268.
- Canoy, M., van der Ploeg, F., & van Ours, J. C. (2006). The economics of books. In V. Ginsburgh, & D. Throsby (Eds.), *Handbook on the economics of art and culture*. Elsevier.
- Caves, R. E. (2000). *Creative industries*. Cambridge, MA: Harvard University Press.
- Chevalier, J., & Goolsbee A. (2003). Price competition online: Amazon versus Barnes and Noble. *Quantitative Marketing and Economics*, 1(2), 203–222.
- Chevalier, J. A., & Mayzlin, D. (2006). The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research*, 43(3), 345–354.
- De Vany, A. S., & Walls, W. D. (1996). Bose-Einstein dynamics and adaptive contracting in the motion picture industry. *The Economic Journal*, 106(439), 1493–1514.
- De Vany, A. S., & Walls, W. D. (2004). Motion picture profit, the stable Paretian hypothesis, and the curse of the superstar. *Journal of Economic Dynamics and Control*, 28(6), 1035–1057.
- Elberse, A., Eliashberg, J., & Leenders, M. (2006). The motion picture industry: Critical issues in practice, current research and new research directions. Forthcoming in *Marketing Science*.
- Franses, P. H. (2002). Testing for residual autocorrelation in growth curve models. *Technological Forecasting & Social Change*, 69(2), 195–204.
- Geroski, P. A. (2000). Models of technology diffusion. *Research Policy*, 29(4–5), 603–625.
- Hall, B. H., & Khan, B. (2003). Adoption of new technology. In D. C. Jones (Ed.), *New economy handbook*. Academic Press.
- Keller, E. B., & Berry, J. (2003). *The influentials*. The Free Press.
- Mahajan, V., Muller, E., & Wind, Y. (Eds.) (2000). *New-product diffusion models*. New York, NY: Springer.
- Moe, W. W., & Fader, P. (2001). Modeling hedonic portfolio products: A joint segmentation analysis of music compact disc sales. *Journal of Marketing Research*, 38(3), 376–385.
- Moul, C. C. (2006). Measuring word of mouth's impact on theatrical movie admissions. Forthcoming in the *Journal of Economics & Management Strategy*.
- Putsis, W. P. J., & Srinivasan, V. (2000). Estimation techniques for macro diffusion models. In Mahajan, Muller and Wind (2000), pp. 263–294.
- Rogers, E. M. (1995). *Diffusion of innovations*. New York: The Free Press.
- Sorensen A. T. (2006). Bestseller lists and product variety. Forthcoming in the *Journal of Industrial Economics*.
- Sorensen A. T., & Rasmussen, S. (2004). Is any publicity good publicity? A note on the impact of book reviews. Working paper.
- Trajtenberg, M., & Yitzhaki, S. (1989). The diffusion of innovations: a methodological reappraisal. *Journal of Business & Economic Statistics*, 7(1), 35–47.

- Van den Bulte, C., & Joshi, Y. V. (2006). New product diffusion with influentials and imitators. Forthcoming in *Marketing Science*.
- Van den Bulte, C., & Lilien, G. L. (1997). Bias and systematic change in the parameter estimates of macro-level diffusion models. *Marketing Science*, 16(4), 338–353.
- Waldfogel, J. (1993). The deadweight loss of Christmas. *American Economic Review*, 83(5), 1328–1336.