Noname manuscript No.
(will be inserted by the editor)

Evaluating the Use of Internet Search Volumes for Time Series Modeling of Sales in the Video Game Industry

Jukka Ruohonen · Sami Hyrynsalmi

the date of receipt and acceptance should be inserted later

Abstract Internet search volumes have been successfully adopted for time series analysis of different phenomena. This empirical paper evaluates the feasibility of search volumes in modeling of weekly video game sales. Building on the theoretical concepts of product life cycle, diffusion, and electronic word-of-mouth advertisement, the empirical analysis concentrates on the hypothesized Granger causality between sales and search volumes. By using a bivariate vector autoregression model with a dataset of nearly a hundred video games, only a few games exhibit such causality to either direction. When correlations are present, these rather occur instantaneously; the current weekly amount of sales tends to mirror the current weekly amount of searches. According to the results, search volumes contribute only a limited additional statistical power for forecasting, however. Besides this statistical limitation, the presented evaluation reveals a number of other limitations for use in practical marketing and advertisement foresight. Internet search volumes continue to provide a valuable empirical instrument, but the value should not be exaggerated for time series modeling of video game sales.

Keywords technology diffusion · foresight · precedence · Google trends · word-of-mouth · EWOM

Jukka Ruohonen
Department of Information Technology, University of Turku, FI-20014 Turun yliopisto, Finland

Sami Hyrynsalmi
Pori Department, Tampere University of Technology, P.O. Box 300, Pori, Finland

E-mail: juanruo@utu.fi
E-mail: sthyry@utu.fi

JEL classification L8 Industry Studies: Services · L82 Entertainment · Media

1 Introduction

In 2008 the market leader in Internet search, Google, Inc., launched a web service that provided estimates for predicting flu trends around the world. Although the service was later discontinued, it was strongly motivated by the promising time series results that were obtained for predicting influenza and other periodic epidemics of infectious diseases (Ginsberg et al 2009; Polgreen et al 2008). While the evaluation work has continued actively in biology, medicine, and health care (Gluskin et al 2014; Proulx et al 2013; Zhang et al 2015; Zhou et al 2013), search engine time series have been increasingly adopted also for many other predictive purposes, covering such different fields as economics (Dimpfl and Jank 2015; Vosen and Schmidt 2011), marketing (Du et al 2015; Rangaswamy et al 2009), ecology (Wilde and Pope 2013), political science (Reilly et al 2012), and, rather tellingly, research of scholarship and the popularization of science (Segev and Baram-Tsabari 2010; Vaughan and Romero-Frías 2014). Although there is some prior work for modeling product sales with relation to Internet search data (Hu et al 2014), thus far, time series applications have been limited in empirical research questions related to the video game industry. This empirical paper sets to the fill the gap.

The larger background is far from being novel, of course. The so-called word-of-mouth (WOM) advertisement was actively researched already in the 1960s for understanding the dynamics related to pioneers, early adopters, latecomers, and non-adopters (Arndt 1967). In fact, in all likelihood, purchasing decisions have been influenced by communication and social dynamics even for centuries. The historical emergence of the Internet brought a new letter E to the WOM abbreviation, nevertheless. Online product reviews, blogs, discussion boards, social media, and related characteristics turned the scholarly attention towards specific forms of electronic word-of-mouth (EWOM) advertisement. Seconded by the fact that search engines continue to dominate the Internet-wide dissemination of information to consumers, this empirical paper approaches mass-scale EWOM effects by evaluating whether Internet search volumes can help at predicting global video game sales volumes. The theoretical background comes from the classical product life cycle literature and the statistically oriented diffusion tradition (Burmester et al 2015; Cox 1967; Dellarocas et al 2009; Chintagunta et al 2009). The evaluation is based on four broad facets.

The paper (1) assesses the underlying theory and assumptions, (2) verifies methods, (3) evaluates data, and (4) validates operational requirements for practical foresight (Piirainen et al 2012). The first two facets are approached with the concept of precedence, which is also known as Granger causality.[1] Namely, (a) the primary empirical question is whether the current weekly volumes of video game sales are affected by the past volumes of Internet searches, or the other way around. The empirical analysis further (b) evaluates the persistence of the time series dynamics in the two volumes, (c) addressing also the effect product launches. Following the recent time series evaluations in economics and marketing (Choi and Varian 2012; Dimpfl and Jank 2015; Du and Kamakura 2012), these three questions are examined with a bivariate vector autoregression (VAR) model, using a longitudinal dataset comprised of nearly a hundred video games. Although a brief forecast experiment is carried out to evaluate the fourth facet, more emphasis is placed on the evaluation of the three questions that can be also motivated theoretically.

With this evaluation, the paper contributes to the domain of video game sales modeling in general, and the subdomain of EWOM research in particular. Besides introducing and evaluating a new and novel source of empirical material for observing EWOM in the video game industry, the results presented allow also contemplating about a few practical implications for marketing and advertisement. In a bigger interdisciplinary picture, the paper further adds the video game industry to the application domains of Internet search volume modeling. The addition is important in its own right. In contrast to the globally occurring but seasonal influenza trends, many variables in the video game industry are directly controllable by the associated companies, which opens the door for theoretically more intriguing time series modeling related to search volumes.

[1] In this paper, the concept of precedence is understood in the sense of Granger's (1988) classical exposition. While the term Granger causality is often used interchangeably, causality itself is a much more problematic concept, of course – as are the related concepts such as exogeneity (here see, e.g., Lütkepohl 2005). This terminological disclaimer should be kept in mind throughout the paper.

2 Background

The theoretical background can be motivated by relying on the so-called diffusion tradition, which has been a prominent interdisciplinary scholarly branch for studying the evolution and propagation of different technological phenomena (e.g., Chintagunta et al 2009; Cortada 2013; Hewing 2011; Ruohonen et al 2015a). For the purposes of this paper, the concept of diffusion itself can be loosely understood as an observable mechanism via which a phenomenon propagates in a population through time, eventually reaching a saturation point and the potential final demise. While this characterization provides a large amount of theoretical anchors, the scope is restricted against the constraints imposed by time series modeling. For framing the theoretical background against these constraints, a few preliminary remarks should be made regarding the approach adopted.

2.1 Approach

A few exploratory statistical assessments are carried out for evaluating whether the results vary across some specific games, but the theoretical focus is still strictly longitudinal. That is, different cross-sectional abstractions (such as gaming platforms, genres, franchises, or geographic markets) are deliberately excluded from the forthcoming theorizing and empirical analysis.

Moreover, only two observable time series quantities are used under a weekly sampling frequency: the volume of video game sales (S) and the volume of Google searches (G). The specific word *volume* can be also used for placing the paper into a specific branch in the extant EWOM literature (cf. Hyrynsalmi et al 2015; King et al 2014). In other words, the fundamental presumption is that the mere amount, the volume of global Internet-wide EWOM, as proxied by G, can help at predicting global video game sales volumes, as measured by S_t. Given this overall expectation, the forthcoming discussion derives six general conjectures for guiding the later empirical time series evaluation. The term conjecture underlines that the knowledge about search volumes is arguably still too incomplete for drawing definite conclusions and contesting strict hypotheses. The same applies to methodological choices. As Choi and Varian (2012) argue, Internet search volume modeling is still an exploratory field, meaning that simple time series models are preferable for better understanding the basic correlations and directions for model refinement.

2.2 Motivation

The use of conjectures does not apply to the larger scholarly background. In particular, there is an ample amount of empirical evidence for the fundamental assertion that electronic word-of-mouth is correlated with product sales (for literature reviews see Dellarocas et al 2009; Hyrynsalmi et al 2015; King et al 2014). Also the electronic commerce research of the video game industry has progressed rapidly in recent years (Chintagunta et al 2009; Hyrynsalmi et al 2016; Kim et al 2014; Landsman and Stremersch 2011; Marchand and Henning-Thurau 2013). Thus, a sufficient theoretical rationale and enough empirical evidence exist for linking search volumes to Internet-wide EWOM, and for motivating the focus on the video game industry.

Even when limiting the attention to the longitudinal dimension and the theoretical diffusion tradition, the video game industry emerges as a particularly interesting case. Consider the vertical and horizontal industry structures as a brief motivating example.

Although game development may be rapid and agile for small games released on mobile platforms, the industry has long been segmented further into personal computer (PC) and console markets. This long-standing horizontal segmentation illustrates the importance of questions related to "*multi-homing*", that is, the development (or porting) of games for different gaming platforms and distribution channels (Hyrynsalmi et al 2016; Landsman and Stremersch 2011). The industry tends to further exhibit a vertical structure; analogous to book publishing, for instance, games are typically developed, published, and distributed by distinct companies (Williams 2002). As the cumulative adoption rates among consumers are usually rapid, typically following characteristically smooth non-linear trends (Chintagunta et al 2009), different diffusion questions are present at different levels of the vertical structures. Due to the horizontal segmentation, the diffusion dynamics vary also across platforms and distribution channels.

For producers and publishers, respectively, concepts such as "*time-to-delivery*" and "*time-to-market*" imply different schedules and deadlines. In fact, the video game industry is notoriously famous for different delays and schedule overruns. For better optimizing the scheduling, a key element is related to better understanding of the consumer behavior. In terms of marketing, it can be crucial for correctly approximating the optimal timing of marketing campaigns, including the increasingly relevant but typically fixed-period online advertisement "banners" (Chen et al 2016). An analogous example would be the commonplace information disclosure in video game events, including the premium Electronic Entertainment Expo. Against this backdrop, it is important to know whether the timing of new releases affects not only sales but also the large-scale EWOM in the Internet. In terms of EWOM itself, the video game industry is particularly relevant due to the active and predominantly Internet-based gaming community (Jöckel et al 2008). The social aspects of gaming underline the importance of opinion leaders, game reviews, different wiki-sites and other community platforms, modification packs developed by volunteers, video streaming services, and numerous other electronic commerce characteristics explicitly or implicitly related to EWOM and the social propagation of information.

For putting these motivating points into a more practical context, consider two recent and novel video games released in 2016, the mobile augmented reality game Pokémon Go and No Man's Sky with its algorithmically generated content. The latter game was widely promoted prior to the actual release, which, however, received largely negative reviews, allowing to question whether the publicity campaigns were correctly optimized. The former game is also highly illustrative. In fact, the release of the new Pokémon game was almost like a stylized textbook example on how to create publicity at a grand scale. However, at the time of writing, it seems that the global social phenomenon that the release was able to generate lasted only for a short period. After release in early July 2016, the adoption peaked in mid-July, and then started to decay (Suber-Jenkins 2016). This rapid trajectory is a good example on the diffusion patterns that are typical to video games.

2.3 Diffusion

Theoretically, it can be asserted that video game sales volumes are driven by the traditional S-shaped diffusion curves, which lead to the classical staged product life cycle models and their corresponding but varying marketing premises (Cox 1967; Polli and Cook 1969). In essence, products are assumed to live through different theoretical stages, such as introduction, momentum, maturity, and the eventual decay. These theoretical stages are understood to reflect empirically observable growth rate periods, such as acceleration, maximum slope, and deceleration, which can be derived from the typical S-shaped curves (Ruohonen et al 2015a).

Then, the current diffusion stage of a game places constraints on the suitable marketing and advertising strategies. The mass of late comers likely needs a different strategy than the relatively small but rapidly adopting amount of initial consumers, and so forth.

Most entertainment products, whether movies, music, or video games, have only a short life cycle; sales start to diminish quickly (Burmester et al 2015; Zhu and Zhang 2010). Also the structure of the video game industry places constraints on the life cycles. A good example would be the continuing polarization into hardware and software segments. While hardware sales may affect software sales, or the other way around (Chintagunta et al 2009; Kim et al 2014; Landsman and Stremersch 2011), hardware platforms and their life cycles ultimately restrict also the life cycle of software products (cf. Marchand 2015). The transition may take long – as exemplified by the contemporary migration to the new generation game consoles – but when a platform has finally reached the stage of demise, the products for the platform are unlikely to generate noteworthy revenue. Although such long-run scenarios are important for understanding the evolution of the video game industry as a whole, the short-run diffusion is more relevant for the sales dynamics of individual games.

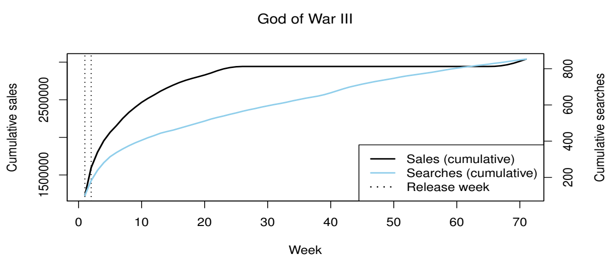

Fig. 1 An illustration of Video Game Sales Dynamics

Because the life cycles are short, the presumed S-shaped trends should show maximum growth rates early on. Indeed, many entertainment industry segments exhibit so-called "*big bang dynamics*"; sales are strongly influenced by the performance during the initial launch period (Dellarocas et al 2009). This effect is best illustrated with concrete empirical data. As a motivating example, consider thus the weekly cumulative sales of a video game called God of War III. By using a dataset that is later described in detail, both the S-shaped trend and the big band dynamics are highly visible in the illustration shown in Fig. 1.[2] While growth was very rapid during the initial few months, in less than a year, a point of diminishing sales was reached. Also the search volumes followed a strikingly similar pattern. Given this general life cycle and diffusion reasoning, the following three Conjectures C_1, C_2, and C_3 seem reasonable theoretically and empirically.

[2] Although cumulative sales and search volumes are used for a few illustrations, formal estimation is carried out with the non-manipulated weekly S_t and G_t time series.

C_1 *In the context of video game industry, volumes of sales and searches tend to exhibit only relatively short time series "memories", meaning that the statistical effect of past volumes on current volumes tends to fade away approximately after a month or two.*
C_2 *The timing of new video game releases affects positively the weekly volume of video game sales.*
C_3 *New releases tend to peak Internet search volumes.*

Also classical (E)WOM presumptions follow. If early adopters drive the rapidly accelerating early growth rates in sales, special care should be taken to lure this segment, which, by hypothesis, propagates the received information to more hesitant consumers. By further considering the whole consumer population as a social network, the propagation of information within the network typically yields a familiar S-shaped curve (De Nooy et al 2011). This reasoning further leads to the famous concept and theory of network effects (Katz and Shapiro 1985), which not only characterize the positive influence of an existing user base upon new adopters, but also extend to the noted industry connection between hardware and software (Marchand and Henning-Thurau 2013). Then, in case the social propagation mechanisms truly are as important as widely believed, it is also relevant to know whether these mechanisms tend to precedence sales peaks, or the other way around.

2.4 The Precedence Puzzle

It is not difficult to demonstrate the empirical plausibility of the correlation between Internet search volumes and video game sales volumes. Thus, consider another motivating illustration in the form of Fig. 2. For this particular game, the two time series, S_t and G_t, have followed even surprisingly similar trajectories. Also the big bang dynamics are again visible; both sales and searches peaked around the shown product launch for a particular geographical market.

The diffusion background allows to question the longitudinal direction of causality, however. It can be argued, for instance, that the EWOM effects are only present during the later life cycle stages (see Zhu and Zhang 2010, who use a four month cut-off point in sales to dichotomize a game's state of popularity). It is also possible to argue with a more fine-grained granularity.

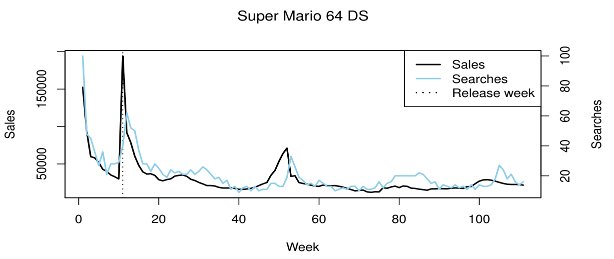

Fig. 2 Another Illustration of Video Game Sales Dynamics

On one hand, the past amount of Internet searches may influence the current amount of sales because of successful marketing and advertisement campaigns. This potential effect of past search volumes would exemplify a prevalent characteristic of the entertainment industry; the commonplace strategy to create sufficient "*buzz*" before product launches via product pre-announcements and related techniques (Hennig-Thurau et al 2015; and more generally Su and Rao 2010). When interpreted against the larger life cycle literature, such buzz fabrics would also signify the so-called pre-release "*shadow diffusion*", which may foreshadow sales volumes due to prior adoption decisions and even concrete pre-orders (here see Burmester et al 2015; and not to be confused with piracy; see Hewing 2011). It is also possible to argue in favor of the reverse – but not necessarily contrary – direction of longitudinal causality.

Thus, on the other hand, past sales may affect, either positively or negatively, the present EWOM volume of a game due to consumers' increased knowledge about the video game (Situmeang et al 2014). For instance, a blockbuster game is likely to receive attention also from the latecomers precisely because of the large amount of past sales. It is also possible that consumers who have already bought such a game will subsequently conduct Internet searches for information about the game – an assumption that has been also known as a "*post-purchase contagion*" effect (Hu et al 2014). This general direction of causality reasoning is also hinted by the illustration in Fig. 2, which shows that the search volumes peaked slightly after the sales volumes had already peaked a week or two earlier.

Finally, keeping in mind the weekly sampling frequency, it is also entirely possible that the current weekly sales and search volumes are "instantaneously correlated". These three different alternatives are formalized in Conjectures C_4, C_5, and C_6, respectively. Unlike with the earlier three conjectures C_1, C_2, and C_3, the following listing carries a more inferential statistical logic.

C_4 *The past volume of weekly Internet searches made for a video game tends to either increase or decrease the current weekly sales volume of the video game.*

C_5 *The past volume of video game sales tends to influence, either positively or negatively, the volume of current Internet searches.*

C_6 *In addition to C_4 and C_5, the current search volume of a game correlates with the current volume of sales; the volumes are "instantaneously correlated".*

The six conjectures presented allow to contemplate about practical foresight value for marketing and advertisement purposes. If C_1 and C_6 are accepted but C_4 is rejected, for instance, it would seem sensible to recommend that the generation of buzz would be worthwhile with short-run campaigns. The conjectures also bind the paper into the larger interdisciplinary research domain that has examined the precedence questions related to Google search data (see, in particular, Dimpfl and Jank 2015; Hu et al 2014). The answers to the six conjectures also enlarge the arguably still limited (Brynjolfsson et al 2010) existing empirical knowledge regarding the impact of Internet search volumes on different real-world phenomena.

3 Data

The dataset covers 96 video games (see Fig. 3). The sample sizes vary: $T_1, T_2, \ldots, T_{96} \in [21, 279]$ weeks. No biases are present towards any specific geographical market, platform, or genre. The sampled games range from big multi-platform Hollywood-style productions, such as Grand Theft Auto V, to smaller games released for handheld devices and smartphones. While the sample size is on the small side for observing the video game industry in general, the dataset is neither atypical (cf. Dellarocas et al 2009) nor ill-prepared for time series analysis. The aspect of quality, rather than quantity, is reflected in the weekly sampling frequency, which allows to observe the short-run dynamics present in the postulated conjectures. Care was also used during the data collection in late January 2016.

3.1 Data Sources

Two sources were used for data collection to obtain the variables of interest. First, release information and global weekly sales statistics were obtained from the business intelligence company VGChartz, Ltd. (2016). In general, the data covers also retain channels, excluding pre-orders. The company obtains these statistics with regular polls from both consumers and retail partners. The polls are accompanied with statistical robustness evaluations and checks against data from other tracking firms and vendors' shipment statistics. The company's sales statistics have been also frequently used in recent empirical research (Kim et al 2014; Marchand 2015; Situmeang et al 2014), which, in general, raises the overall confidence over the validity and reliability of the data for research use.

The observed dataset was not constructed by random sampling of games. After all games with weekly sales statistics were retrieved, the sample was then reduced according to suitability for time series analysis. A large number of games had to be excluded due to insufficient number of weekly sales statistics. A minimum of twenty weekly observations was used as the criterion for inclusion in the sample. This choice is close to the absolute minimum that is required for time series analysis in general.

Second, the standard online trend service of Google, Inc. (2016) was used to obtain the weekly search volumes. Rather than seeking for generality (cf. Vaughan and Romero-Frías 2014), less ambiguous data (cf. Drake et al 2012) was sought by querying each game from the Google's trend service with the exact same name provided by VGChartz, Ltd. (2016), further enclosing the names with quotation marks.[3] This choice also allows reproducing the dataset.

As the obtained sales data refers to worldwide purchasing statistics, also the search trends were specified with a global scope. Again a number of games had to be excluded due to the unavailability of search data, which is explained by the insufficient global search volume. To maintain the weekly sampling frequency, also those games were excluded for which Google provided only monthly statistics. Finally, a few games had to be further excluded because of their names, which were too general to be adequately represented by search engine trends. These exclusions are subjective but sensible, including such games as Cars, Grid, Game of Thrones, Infected, Skate, Robots, Gun, Pure, Rock Band, Catherine, Hannah Montana, Madagascar, NBA, and Gretzky NHL. Already a visual evaluation reveals that the search trends for these few games cannot have been fully generated by people specifically searching for these

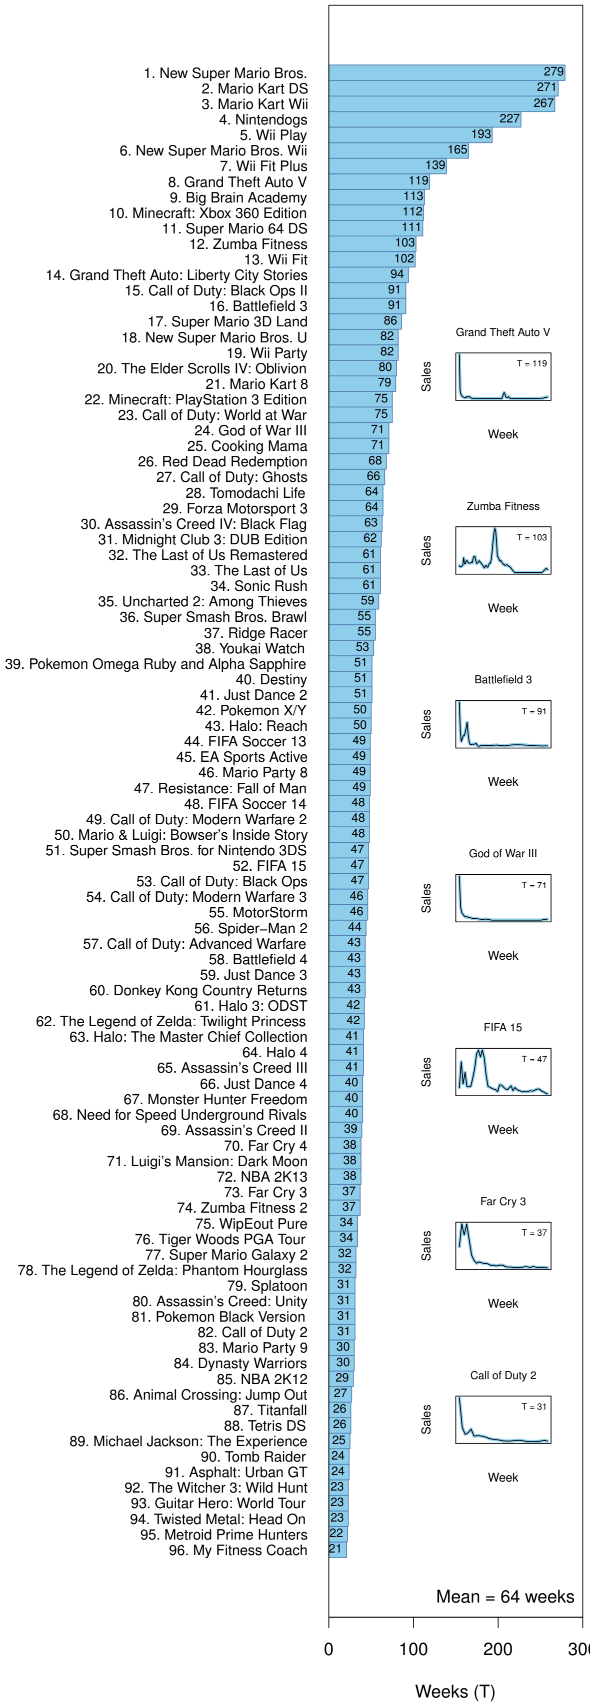

Fig. 3 The Dataset

[3] A few details are also worth remarking. First, Google defines some characters (such as the plus sign) invalid; these were omitted from a few names. Second, release information (see Section 3.2.2) was also queried with the same strategy, and, hence, special editions and other non-standard release types are counted as separate games insofar these carry different names. Third, the platform-specific releases of a few outliers (notably, Super Mario Bros.) are classified by VGChartz, Ltd. (2016) as separate games, although these are collapsed into one composite for most other sampled games.

broadly entitled video games, especially when further reflected against the release dates and sales trends.

3.2 Operationalization

The dataset contains two continuous time series variables and two deterministic dummy variables. A few remarks should be made about the operationalization of both variable types.

3.2.1 Continuous Variables

The data collection procedure ensured full synchronization between S_t and G_t. In other words, the two data sources share the same calendar time continuum; the seven day sampling frequency starts from the 26th of December 2004 and ends to the 26th of December 2015 for both sources. Unlike with some related research problems (Segev and Baram-Tsabari 2010; Vosen and Schmidt 2011), there is no need for calendar time aggregation into monthly or annual frequencies. Yet, there exists a tricky problem of missing values in the weekly sales statistics of many sampled games. That is, an imaginary vector of raw sales might look like

$$[\mathrm{NA}, \ldots, \mathrm{NA}, \ldots, \mathrm{S}_i, \mathrm{NA}, \mathrm{S}_{i+2}, \mathrm{S}_{i+3}, \mathrm{NA}, \ldots, \mathrm{NA}]', \quad (1)$$

where the symbol NA denotes a missing value. Plain dropping of the missing values is not a plausible option because the resulting time series would be irregular in calendar time.

Rather than excluding observations, the problem is solved, for each game separately, (a) by first restricting the estimation period to the period between the first and the last available non-missing weekly time point. This reduction would result $[\mathrm{S}_i, \mathrm{NA}, \mathrm{S}_{i+2}, \mathrm{S}_{i+3}]'$ for the example vector (1). Then, (b) the remaining missing values are interpolated using a cubic spline approximation (for splines and interpolation see, e.g, Hastie et al 2011). If $\mathrm{S}_i = 1$ and $\mathrm{S}_{i+2} = \mathrm{S}_{i+3} = 3$ in the example vector, a value 2.333 would be used for the single missing value, for instance. It was also necessary to (c) exclude the games that required more than a year of interpolation approximations. Finally, keeping in mind the supposedly non-linear diffusion dynamics of the weekly sales volumes, (d) in case the cubic spline technique resulted negative values, these were replaced with zeros because negative sales are impossible.

While the exclusion threshold of 52 interpolated observations is arbitrary in the case (c), it is arguably sufficient to ensure decent approximations. It can be also remarked that without the criterion, there would be prohibitively long periods for interpolation. For instance, VGChartz, Ltd. (2016) provides a period from circa November 2006 to May 2007 for the game Gears of War. This period contains three missing values, which are all robustly enough interpolated by the cubic spline technique. Yet, the source provides a further small period of data that dates to circa September 2015. Clearly, interpolating the long gap in-between would be too long for the present purposes.

Finally, Google, Inc. (2016) only provides indexed search volumes at the range from a zero to hundred, the maximum value corresponding with the highest relative number of searches in an unknown time interval. As is well-understood, this indexing makes it impossible to deduce about the actual absolute amount of searches (Curme et al 2014; Reilly et al 2012; Segev and Baram-Tsabari 2010; Vaughan and Romero-Frías 2014). The maximum value 100 might refer to a few thousand searches, or it might denote a few hundred million searches – or even more. This general limitation should be kept in mind when interpreting the empirical results, although it does not affect the statistical computations as such.

3.2.2 Deterministic Variables

In general, the construction of suitable release proxies is difficult because video games are typically released for different geographical markets at different times, and these release timings often further vary according to video game platforms (Chintagunta et al 2010). Sometimes there may even be a decade in-between the initial releases and the release of later versions ported to new video game platforms. For instance, Ridge Racer was released for Playstation in 1994, which was much later, in 2004, followed by a release for Playstation Portable.

Analogous historical multi-platform scenarios are approximated constructing three dummy-variables of the following form:

$$d_{it} = \begin{cases} 1 & \text{for any platform-specific release at } t \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

given $i \in \{\mathrm{U.S.}, \mathrm{Europe}, \mathrm{Japan}\}$. Because the resulting vectors may be identical for some games due to the same release dates across the three continents, the three variables were further collapsed into a single

$$d_t = \max\{d_{\mathrm{U.S.},t}, d_{\mathrm{Europe},t}, d_{\mathrm{Japan},t}\}. \quad (3)$$

Holiday seasons are important for video game sales in all three continents (Harada 2007; also Chintagunta et al 2009; Kim et al 2014). To proxy the supposedly most important holiday season for video games sales, the Christmas week is controlled with

$$c_t = \begin{cases} 1 & \text{if the 24th of December is at week } t \\ 0 & \text{otherwise} \end{cases}. \quad (4)$$

It should be noted that c_t likely proxies only a small portion of Christmas sales. While longer effects have been considered (for instance, Marchand 2015 defines $c_t = 1$ for a period between October and December), the operationalization in (4) is justified in the present context due to the relatively short length of some observed time series (see Fig. 3). A further remark should be made also about the geographic dimension.

In theory, it would be relevant to model particularly the EWOM dynamics through the geographical diffusion from a continent to another. For instance, a new release in the U.S. likely increases EWOM volume in the North America, which likely also increases electronic word-of-mouth volume in Europe even when the corresponding European release would not have yet occurred. Nevertheless, the use of the collapsed (3) is justified because of the global scope of both S_t and G_t, and because the geographical release timings have often been outside of the observed period due to the unavailability of data even after the interpolation. The autoregressive estimation strategy also reduces the sample sizes, which may further push the Christmas week or the global release dates outside of a given estimation sample. Moreover, an analysis with specific geographic markets faces additional data collection challenges, which can be at least partially remedied by operating at the global scale.

4 Methods

Estimation is carried out for each game separately by using a bivariate VAR model. Computation is carried out with an R implementation (Pfaff 2008), which follows closely the exposition and implementation (JMulTi) of Lütkepohl (2005) and associates. The forthcoming discussion outlines the basic methodology.

4.1 Model

The primary estimation methodology is based on a bivariate vector autoregression model, possibly augmented with the two deterministic variables. The estimated pth order vector autoregressions are algebraically simple:

$$\mathbf{x}_t = \mathbf{u} + \sum_{i=1}^{p} \mathbf{A}_i \mathbf{x}_{t-i} + \mathbf{b}_1 \mathbf{c}_t + \mathbf{b}_2 \mathbf{d}_t + \mathbf{e}_t, \tag{5}$$

where $\mathbf{x}_t$ is a two-dimensional vector comprised of the two observed (endogenous) time series, $\mathbf{u}$ is a vector of intercepts, $\mathbf{A}_1, \ldots, \mathbf{A}_p$ are coefficient matrices, $\mathbf{b}_1$ and $\mathbf{b}_2$ are coefficient vectors for the two (exogenous) deterministic variables in (4) and (3), respectively, and $\mathbf{e}_t$ is a residual vector. The model is a "vector regression" because there are two dependent (endogenous) variables with their own "mirror image equations". The model is also autoregressive because both of these equations contain lagged values of the dependent variables on the right-hand sides. If p is moderate, the recent past explains the near future.

It should be emphasized that the representation in (5) is a little misleading because either $\mathbf{b}_1 \mathbf{c}_t$ or the term

$$\mathbf{b}_2 \mathbf{d}_t = \begin{bmatrix} \beta_{1,2} \\ \beta_{2,2} \end{bmatrix} \begin{bmatrix} d_t \\ d_t \end{bmatrix} \tag{6}$$

is excluded in case the two dummy variables record no values besides zero. Moreover, if $\mathbf{c}_t$ equals $\mathbf{d}_t$ for all t, only one of the variables is included. As this multicollinearity case occurs for one game, the one included variable proxies the combined release and Christmas effects in this case.

In general, the residual vector is assumed to contain zero-mean error processes that are independent and normally distributed with a positive-definite variance-covariance matrix $\mathbf{V}$. In other words, $\mathrm{E}(\mathbf{e}_t \mathbf{e}_t') = \mathbf{V}$ and $\mathbf{e}_t \sim \mathrm{N}(\mathbf{0}, \mathbf{V})$, where $\mathbf{0}$ is a vector of zeros As such, there is also nothing unusual in the basic time series premises: (a) the model is linear in the parameters that are (b) constant over time, while the (c) error terms are normally distributed (d) without serious time dependencies (Hendry and Juselius 2001). These are all conventional time series premises translated to a multivariate context. In general, as a prior-analysis expectation, all but (c) seem reasonable assumptions; non-normality is to be expected due to the presumed non-linear diffusion dynamics.

4.2 Precedence

To see more clearly why the model is called a vector autoregression, consider the simplest case with $p = 1$, such that only the first lags, S_{t-1} and G_{t-1}, are used on the right-hand sides. By further omitting the deterministic variables for brevity, unpacking the matrix notation yields

$$\begin{bmatrix} \mathrm{S}_t \\ \mathrm{G}_t \end{bmatrix} = \begin{bmatrix} \mu_{\mathrm{S}} \\ \mu_{\mathrm{G}} \end{bmatrix} + \begin{bmatrix} \alpha_{11,1} & \alpha_{12,1} \\ \alpha_{21,1} & \alpha_{22,1} \end{bmatrix} \begin{bmatrix} \mathrm{S}_{t-1} \\ \mathrm{G}_{t-1} \end{bmatrix} + \mathbf{e}_t, \tag{7}$$

where $[\mu_{\mathrm{S}}, \mu_{\mathrm{G}}]'$ is the intercept vector $\mathbf{u}$ in (5) and

$$\mathbf{e}_t = \begin{bmatrix} \varepsilon_{\mathrm{S},t} \\ \varepsilon_{\mathrm{G},t} \end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \upsilon_{11} & \upsilon_{12} \\ \upsilon_{21} & \upsilon_{22} \end{bmatrix} \right) \tag{8}$$

with positive variances ($\upsilon_{11} > 0$ and $\upsilon_{22} > 0$). If the second right-hand term in (7) is multiplied out, that is,

$$\mathbf{A}_1 \mathbf{x}_{t-1} = \begin{bmatrix} \alpha_{11,1}\, \mathrm{S}_{t-1} + \alpha_{12,1}\, \mathrm{G}_{t-1} \\ \alpha_{21,1}\, \mathrm{S}_{t-1} + \alpha_{22,1}\, \mathrm{G}_{t-1} \end{bmatrix}, \tag{9}$$

it is better seen that the lagged value G_{t-1} does not influence S_t when $\alpha_{12,1} = 0$, and the other way around in case $\alpha_{21,1} = 0$. The same generalizes to higher orders. In fact, Granger causality has a simple intuitive meaning in the context. If $\alpha_{12,1} = \alpha_{12,2} = \cdots = \alpha_{12,p} = 0$ for the VAR(p) process in (5), the second series does not Granger cause the first series (Lütkepohl and Reimers 1992). That is to say, each $\mathbf{A}_1, \ldots, \mathbf{A}_p$ would be lower-triangular, the upper-right corner being zeros.

This longitudinal causality reasoning provides means for evaluating C_4 and C_5. Testing is carried out with conventional F-tests by comparing the corresponding unrestricted and restricted models. The third Conjecture C_6, in turn, can be examined by evaluating whether $\upsilon_{12} = \upsilon_{21} = 0$ in (8), and analogously for higher order processes. Thus, in essence, the residual vectors $\varepsilon_{\mathrm{S},t}$ and $\varepsilon_{\mathrm{G},t}$ must be uncorrelated for no instantaneous effects between the variables (Lütkepohl 2005). A thousand bootstrapped computations are further used to derive the associated p-values by using the routine provided in the implementation (Pfaff 2008). Finally, it should be emphasized that the bivariate vector autoregression model is generally estimated equation-by-equation with ordinary least squares. By implication, unless either C_4 or C_5 is rejected, or both are, the bivariate model is not identified in the sense that the two series could be uniquely represented in a single system of equations.

4.3 Stability

The foremost estimation requirement is the stability of the autoregressive process. The most straightforward violation occurs when some of the included time series are not stationary, meaning that their mean and variance do not remain constant over the observed period.

To test the stationarity requirement, univariate per-series testing is possible with the conventional applied tools (Ruohonen et al 2015b), which have been also used to assess stationarity requirements in Internet search volume modeling (Zhang et al 2015). The multivariate VAR context permits also somewhat easier – or at least more "ergonomic" – stationary testing (Juselius 2006). A commonly used testing technique is based on a matrix decomposition in the form of a so-called companion matrix (for technical details see Bårdsen et al 2005; Juselius 2006; Lütkepohl 2005). This technique is also sufficient for the present purposes: the stability of the models is evaluated by observing the eigenvalues of the companion matrix. Then, in essence: if all observed series are stationary, a VAR(p) process is stable, and all eigenvalues are inside or on the unit circle. This stability requirement is easy to evaluate visually by simply plotting the associated eigenvalues, subsequently modifying the plot by adding a circle with a radius of one.

4.4 Order Selection

Determining the optimal order can be tricky. Besides the theoretical problems associated with interpreting C_1 against the integer p in (5), the order selection affects most of the statistical assumptions, including the stability requirement. Moreover, the order selection question exposes a problem of overparametrization; a large p leads to a large number of parameters to be estimated. As can be deduced from (9), the number of estimated parameters can escalate also in small bivariate vector autoregressions. Parsimonious small models are also desirable for interpretation (Dellarocas et al 2009). Fortunately for applied work, there are numerous statistical ways to determine the order of a VAR(p) process.

It is possible to consider, for instance, theoretical plausibility, sequential tests of parameter restrictions, forecast performance, cointegration tests, comparisons with simulations, tests for remaining autocorrelation, and different information criteria measures (Athanasopoulos et al 2011; Lütkepohl 2005). Although the sole use of information criteria metrics has been observed to perform poorly (Athanasopoulos et al 2011), the order selection is done automatically according to a routine (Pfaff 2008) that uses the Bayesian information criterion (BIC) for selecting the order. That is, in essence, the order p corresponds to the VAR(p) model that has the smallest BIC. The rationale to prefer this BIC selection is practical: fine-grained decision making is difficult because it is not plausible to carefully (manually) evaluate each estimated model separately. It should be also remarked that the orders were determined prior to the inclusion of the two deterministic variables. As noted also in Section 3.2.2, this choice is necessary because a large selected order may result a zero-valued $\mathbf{c}_t$ or $\mathbf{d}_t$.

4.5 Diagnostics

Numerous diagnostic checks are frequently used in applied VAR modeling (Bårdsen et al 2005). Two statistical tests are sufficient for the purposes of this paper. Namely, the absence of remaining autocorrelation is evaluated with the so-called Portmanteau test, while the Jarque-Bera test is used to assess the normality assumption. Both are multivariate generalizations of the original tests (Pfaff 2008; for details see Lütkepohl 2005). That is, the latter evaluates whether $\mathbf{e}_t$ is normally distributed, whereas the former is used for assessing that the residuals are not correlated with the past residuals. A brief forecast experiment is also carried out to reassess the results from the precedence tests.

The within-sample forecasting experiment is done by predicting the month-ahead weekly sales, using the last four weeks in a given series as the forecasting target. Rather than seeking good forecasting performance in general, the interest is to compare the predictions with unrestricted VAR(p) models to the predictions done with models in which C_4 is restricted to false. In other words, the familiar restriction $\alpha_{12,1} = \cdots = \alpha_{12,p} = 0$ is approached from a more practical foresight perspective; when the performance improves, $\text{G}_{t-1}, \ldots, \text{G}_{t-p}$ are arguably worth retaining, irrespective of statistical significance. This said, the evaluation of forecasting performance is often no easier than evaluation based on statistical significance. To compare the two types of models, the mean absolute error (MAE) can be used; in the present setting, MAE is defined as the monthly average of the absolute difference between the predicted and actual sales in a four week period. While the metric carries limitations, as does inference with statistical significance, it is sufficient for the present purposes.

5 Results

The results are presented in two steps: model specification precedes the evaluation of the six prior conjectures of interest. The actual evaluation is carried out by focusing on the results that can be said to characterize most of the games in the dataset.

5.1 Specifications

The formal estimation can be started by examining the order selection with the Bayesian information criterion. Because the sample sizes vary from a game to another, the selection can be approached by restricting the maximum process order to five and ten. According to the summary in Table 1, this restriction is unnecessary for most games, however. In other words, a relatively short order of $p \leq 4$ is selected for many games, which translates to a month of past information or less. For a few games, however, even a VAR(10) process is indicted by the BIC selection, which, in general, signals that the bivariate processes do not fit well to all games. The foremost suspect is stationarity.

Surprisingly, however, the supposedly non-linear dynamics in the strongly decaying sales volumes, in particular, do not translate into notable stability problems in the bivariate processes. A possible explanation relates to the large early sales amounts (see Table 2). In other words, the big bang dynamics are highly visible: the early weeks are crucial for the overall sales performance; on average across games, the first three weeks amount to about twenty percent of all sales in the sample. These short early fluctuations may not be enough for forceful non-stationarity to emerge, especially since the later growth rates generally decay in a rather stable manner with small increments.

Table 1 VAR(p) Orders (BIC selection, no. of games)

	Order									
	1	2	3	4	5	6	7	8	9	10
Max $p = 5$	33	27	9	8	19	–	–	–	–	–
Max $p = 10$	26	21	6	3	3	5	6	3	8	15

Table 2 Share of Early Sales (% of total absolute sales)

	Min	Max	Mean	Std. Dev.
The first week (S_1)	0.6	45.5	11.2	11.6
The second week (S_2)	0.2	19.4	6.0	4.4
The third week (S_3)	0.0	13.9	4.9	3.5

The stability observation is illustrated in Fig. 4 by plotting the (complex) eigenvalues of the companion matrices under three different specifications. The best specification is at the middle. That is, only one eigenvalue is outside of the unit circle.

The first two plots, (a) and (b), refer to the fitted VAR($p \leq 5$) models with and without (4), respectively. It is evident that the Christmas holiday dummy variable slightly increases the instability of the processes. A likely statistical reason relates to the fact that many games are deliberately released around Christmas. When the deterministic Christmas effect is omitted, which seems reasonable, the $p \leq 5$ restriction results stable processes for all but one outlying game, Tetris DS. The modulus of the associated eigenvalue is 1.074 for this outlier, which is problematic but not overly large. While there are also a few borderline cases in terms of stationarity (cf. Juselius 2006), the majority of the points remain rather neatly within the unit circle. The same conclusion does not apply when (4) is omitted but the maximum order is raised to ten, as clearly seen from the last plot (c) in Fig. 4. Thus, the remaining evaluation is carried out by excluding the Christmas effects for all 96 bivariate VAR($p \leq 5$) processes.

The $p \leq 5$ restriction contributes to the presence of some remaining autocorrelation in some of the estimated models, however. This can be seen from the summary in Table 3, which tabulates the relative percentage share of games for which the two discussed di-

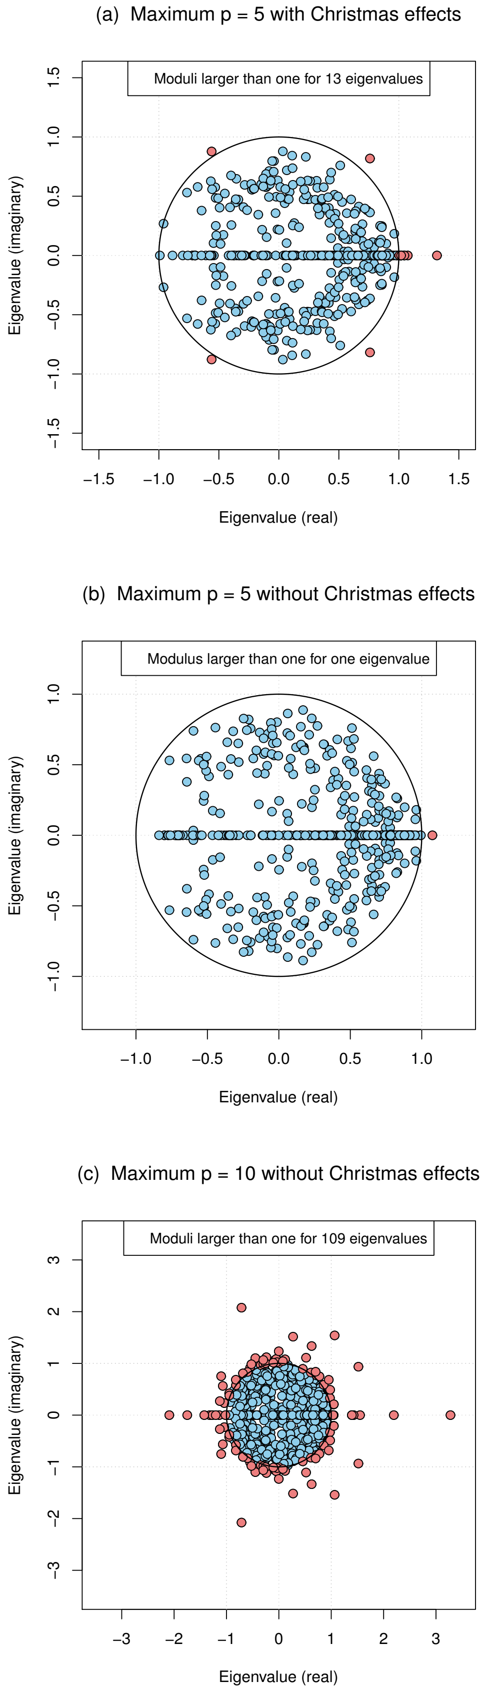

Fig. 4 Eigenvalues of three VAR specifications

Table 3 Diagnostics Checks (% of rejections at 5 % level)

AR($p+1$)	AR($p+5$)	AR($p+10$)	Normality
78.1 %	56.2 %	28.1 %	78.1 %

Specification: VAR($p \leq 5$), (4) omitted

agnostic checks (see Section 4.5) resulted p-values less than 0.05. That is, the table shows the relative share of rejections for the null hypotheses of no remaining autocorrelation and multivariate normality. As expected, the latter is also problematic; only about 22 % of the models accept the null hypothesis of normality. Nevertheless, it is also worth emphasizing that the remaining autocorrelation effects decline rather fast. While for 75 games the first lag after the VAR order is problematic, only a few models would benefit from the generally unstable high-order VAR($p \leq 10$) specification.

In general, the autocorrelation and other potential statistical problems presumably reduce the forecast accuracy for many games. Because the stability requirement is satisfied with the preferred specification (put aside the single outlier out of 79), sufficient fits are available for trying to answer to the six conjectures of theoretical interest.

5.2 Evaluation

The BIC-based selection tentatively supports the first Conjecture C_1 even with the longer $p \leq 10$ restriction (see Table 1). The illustration in Fig. 5 allows to also conclude that Conjectures C_2 and C_3 attain sufficient but approximate support. The shown information refers to the per-game estimates for the $\beta_{1,2}$ and $\beta_{2,2}$ coefficients from (6). For the 60 bivariate models to which (3) could be included, most of the coefficients have positive signs. The magnitudes are also relatively high particularly for the Internet search volumes. Thus, when compared to the normal weeks without new game releases, the release weeks tend to generally exhibit higher sales volumes and larger amounts of Internet searches.

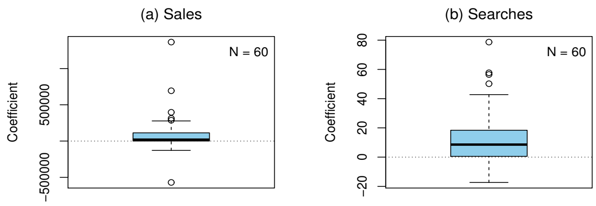

Fig. 5 Release Effects ($p \leq 5$ without Eq. 4)

The two precedence assumptions fail to realize. A half of the decision making material is summarized in

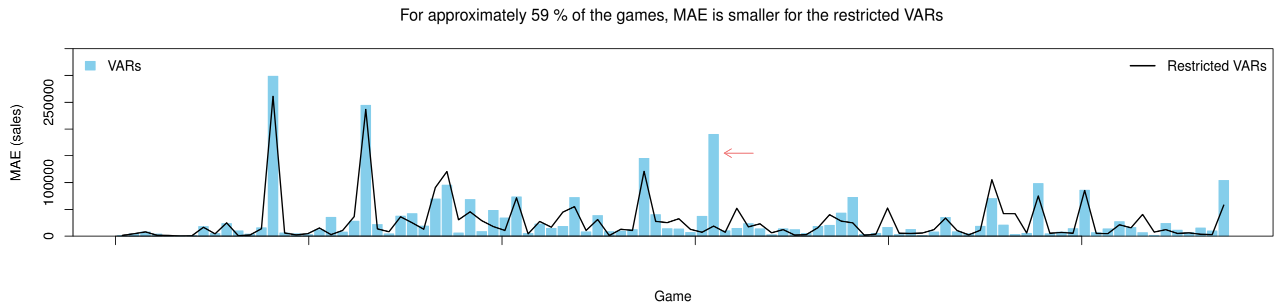

Fig. 6 Sales Forecasts with Restricted and Unrestricted VAR($p \leq 5$) Models (month-ahead, MAEs)

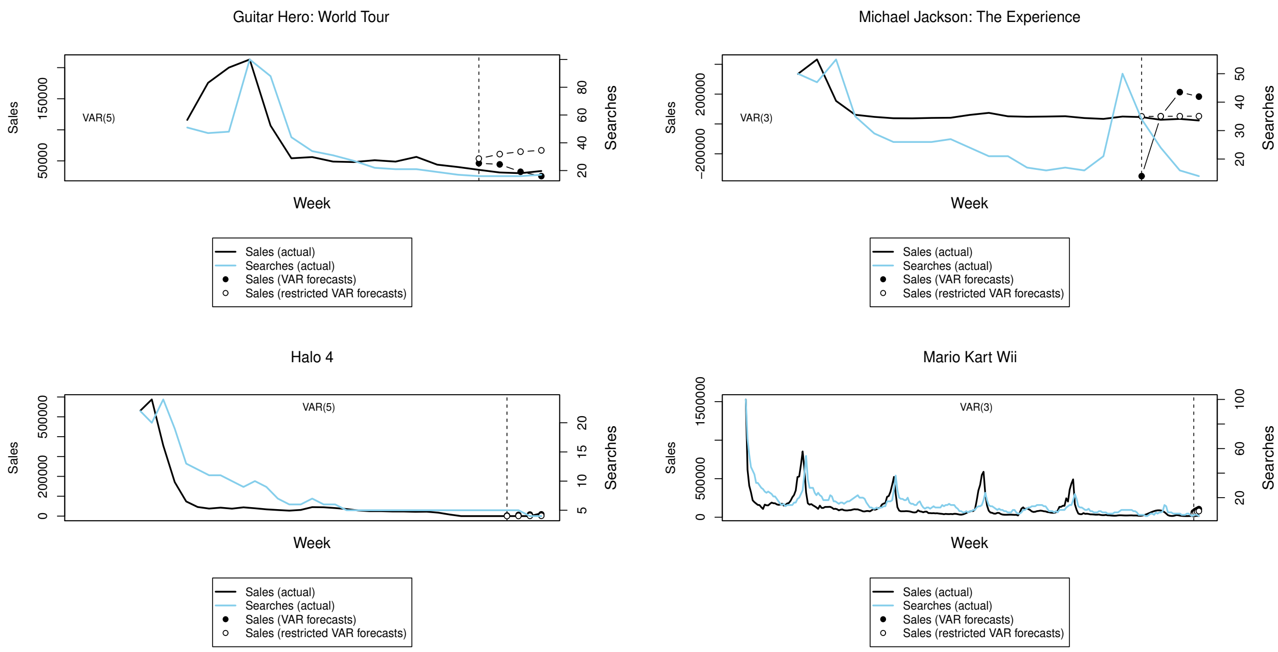

Fig. 7 Four Sales Forecasts with Restricted and Unrestricted VAR($p \leq 5$) Models (month-ahead)

Table 4 Precedence Tests (% of rejections at 5 percent level)

Conjecture	%
C_4 vs. $H_0 : \alpha_{12,1} = \cdots = \alpha_{12,p} = 0$	20.8 %
C_5 vs. $H_0 : \alpha_{21,1} = \cdots = \alpha_{21,p} = 0$	13.5 %
C_6 vs. H_0: no instantaneous effect	45.8 %

Specification: VAR($p \leq 5$), (4) omitted

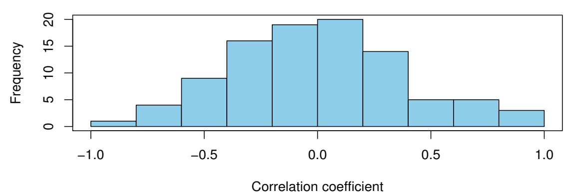

Fig. 8 Residual Correlations ($p \leq 5$ without Eq. 4)

Table 4. According to the numbers on the first two rows, on average, Granger causality between the two variables cannot be used to generally characterize the whole dataset. Only for a minority (21 %) of games, $G_{t-1}, \ldots, G_{t-p}$ influences S_t according to the reported F-tests. The reverse applies even for a smaller minority.

On the other hand, the instantaneous effects between S_t on G_t apply to about 46 % of the sampled games. These instantaneous effects are shown also in the lengthy range $[-0.95, 0.93]$ for the correlation coefficients between the residual series $\varepsilon_{S,t}$ and $\varepsilon_{G,t}$. From the 96 per-game correlations shown in Fig. 8, about 48 % are significant at 5 % level. The formal test results in Table 4 are in accordance with the amount.

A brief evaluation of the forecasting performance supports these observations. The other half of the material can be thus presented in the form of Fig. 6. The foremost conclusion is the small divergence between the mean absolute errors from the full VAR($p \leq 5$) models without Christmas effects and the corresponding restricted models without $\mathrm{G}_{t-1}, \ldots, \mathrm{G}_{t-p}$ on the right-hand sides. While the restricted models exhibit slightly better performance, there are also many exceptions to both directions. For instance, one visible outlier (marked with an arrow in Fig. 6), shows much worse performance with the full unrestricted specification. This small detail is further illustrated in the upper-right plot of Fig. 7, which indicates that a short-run fluctuation in the search volumes affected the sales forecasts, although the actual sales continued to follow the stable path in the month-ahead forecast window. The upper-left side plot exemplifies a contrary case.

All in all, only modest performance gains are available with the additional information from search volumes. This said, it should be emphasized that not all of the forecasts are bad as such. As exemplified by the bottom-right plot in Fig. 7, for some games, there are visible instantaneous effects between the two volume series, and these can also reduce the forecast mean absolute errors. These observations lead to a further question about whether the generally weak support for the precedence puzzle is specific to some particular games.

5.3 Assessments

A few assessments can be computed for a brief evaluation of whether the sample characteristics vary according to the results for the three Conjectures C_4, C_5, and C_6 for which formal statistical tests are available. By classifying the games to binary categories according to whether these pass the precedence tests (see Table 4), it is possible to quickly evaluate whether the per-game time series lengths and absolute sales amounts vary. The results from the computation are shown in Table 5.

The numbers on the first row indicate that there are some differences between the length of the time series particularly in terms of Conjecture C_5. That is, the median time series length (T) is higher among those games for which past sales volumes tend to affect the current search volumes. Interestingly, the contrary does not apply as strongly. When turning to the second row, it is also evident that absolute sales amounts are higher for games that support C_5 and C_6. Although it should be kept in mind that the per-game absolute sales amounts vary according to T, these observations might be used as weak evidence for justifying a further assertion that particularly sales volumes of well-selling or blockbuster games may affect Internet search volumes, but not necessarily the other way around.

There seem to be no clear cross-sectional groupings behind the results, however. Consider thus Fig. 9, which shows the about 21 and 14 percent of games that support C_4 and C_5, respectively (see also Table 4). For instance, on the left-hand side, there are games for mobile and console platforms, relatively old games (such as MotorStorm and Twisted Metal: Head On) and more recent games (including Witcher 3 and Far Cry 4). In particular, the results do not seem to vary systematically across the observed per-game calendar time periods. The observation is interesting because search engine technologies and people's use of search engines have both evolved from the mid-2000s when the observed Twisted Metal game was released.

On both sides, there are also games that represent wildly different genres, ranging from racing games and so-called first-person shooters to the music and fitness genres. The situation becomes more complex by making a further breakdown in terms of the support for the instantaneous effects (C_6). However, it is again easy to pinpoint different contradictions between the two listings in Fig. 10. For instance, the results are not uniform across some game franchises, including Far Cry 3 and 4 as well as Battlefield 3 and 4. As with the earlier precedence results (see Fig. 9), no uniform pattern can be observed in terms of the per-game calendar time periods. Although a more thorough statistical analysis would be required for a definite answer, it can be summarized that conventional cross-sectional abstraction may not be enough for explaining the diverging results across the observed video games.

6 Discussion

This empirical paper evaluated the potential of Internet search volumes for time series modeling of video game sales. By framing the forthcoming summary to the four specified evaluation facets related to theory, methods, data, and practice (Piirainen et al 2012), the remainder of this paper first summarizes the key empirical results against the specified six conjectures and the existing literature. The summary is followed by a few brief notes about managerial implications and research limitations.

6.1 Summary of Results

The results are summarized in Table 6. Most importantly, the dual precedence assertions, C_4 and C_5, are both rejected for the majority of games. Search volumes do not seem to generally Granger-cause video game sales volumes. The reverse applies to even fewer games.

Table 5 Sample Characteristics According to Conjectures

	G_t Granger-causes S_t		S_t Granger-causes G_t		Instantaneous effects	
	Yes (C_4)	No (vs. C_4)	Yes (C_5)	No (vs. C_5)	Yes (C_6)	No (vs. C_6)
Median T	42.5	47.5	54.0	44.0	47.5	44.0
Median $\sum_{t=1}^{T} S_t$	2834164.1	2936695.4	3919701.8	2772403.9	3603660.2	2748239.0

Specification: VAR($p \leq 5$), (4) omitted

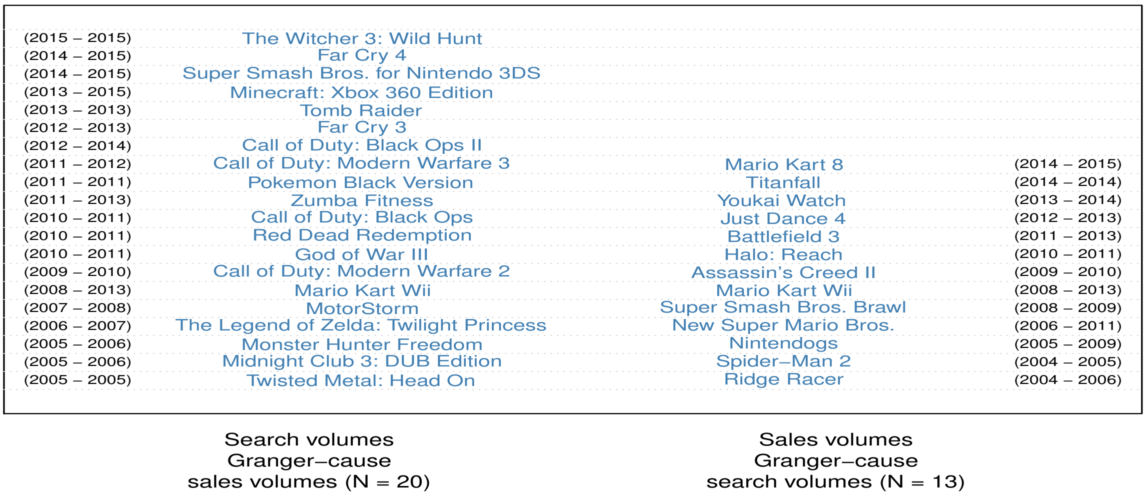

Fig. 9 Breakdown by Precedence Effects (observed calendar time periods in parenthesis on left and right of the groupings)

At first glance, the conclusion may appear as somewhat disappointing for the question of using Google search data for modeling sales volumes of video games in general. Because C_6 can be partially accepted – there are often instantaneous effects – this line of thought is slightly unwarranted, however. For many games, the two series are strongly correlated, following rather similar short-run trends, but these correlations largely occur contemporaneously. Therefore, when predicting the weekly sales volumes in the video game industry, the results rather bespeak for a simple univariate autoregressive model that utilizes the current search volumes as an exogenous right-hand side variable.[4] By possibly further augmenting such a model with moving averages or other time series properties, a classical time series path would be readily available for further improvements (for related models see Du and Kamakura 2012). Model benchmarking is also one plausible option for further empirical research.

While the generalizability of the results to other application domains is limited, some of the findings are in line with the existing literature. In particular, the estimated models support the observation that lagged sales may have only a limited impact current search volumes (Hu et al 2014). There are also some signals in the literature for instantaneous effects between domain-specific time series and Google search volumes. For instance, lagged search volumes have been considered but still ultimately omitted, often without explicit evaluation (Choi and Varian 2012; Ginsberg et al 2009). Although instantaneous effects have been observed previously also under a comparable time series setup, it is apparent that also past search volumes may improve predictions in some applications (Dimpfl and Jank 2015). Thus, it is important to emphasize that the results apply only to the observed video games, and, moreover, there are also many exceptions to this conclusion.

In fact, over a half of the games seem to exhibit no forceful instantaneous effects either. What is more, forecasting performance of sales with the VAR mod-

[4] To clarify, a simple leading indicator model (Choi and Varian 2012), $S_t = \mu + \sum_{i=1}^{p} \alpha_i S_{t-i} + \beta_1 G_t + \beta_2 d_t + \varepsilon_t$ could be used as a simple benchmark model for the VAR models, among other algebraic variations of the equation. Another option might be to investigate the effect of deterministic trend functions in the VAR specifications. Such models should be further benchmarked against the common (Chintagunta et al 2009; Ruohonen et al 2015a) non-linear diffusion models. Finally, it should be remarked that predicting a game's later life cycle may not be a realistic scenario in practice. Therefore, further empirical experiments are required for forecasting time periods that are closer to video game release dates.

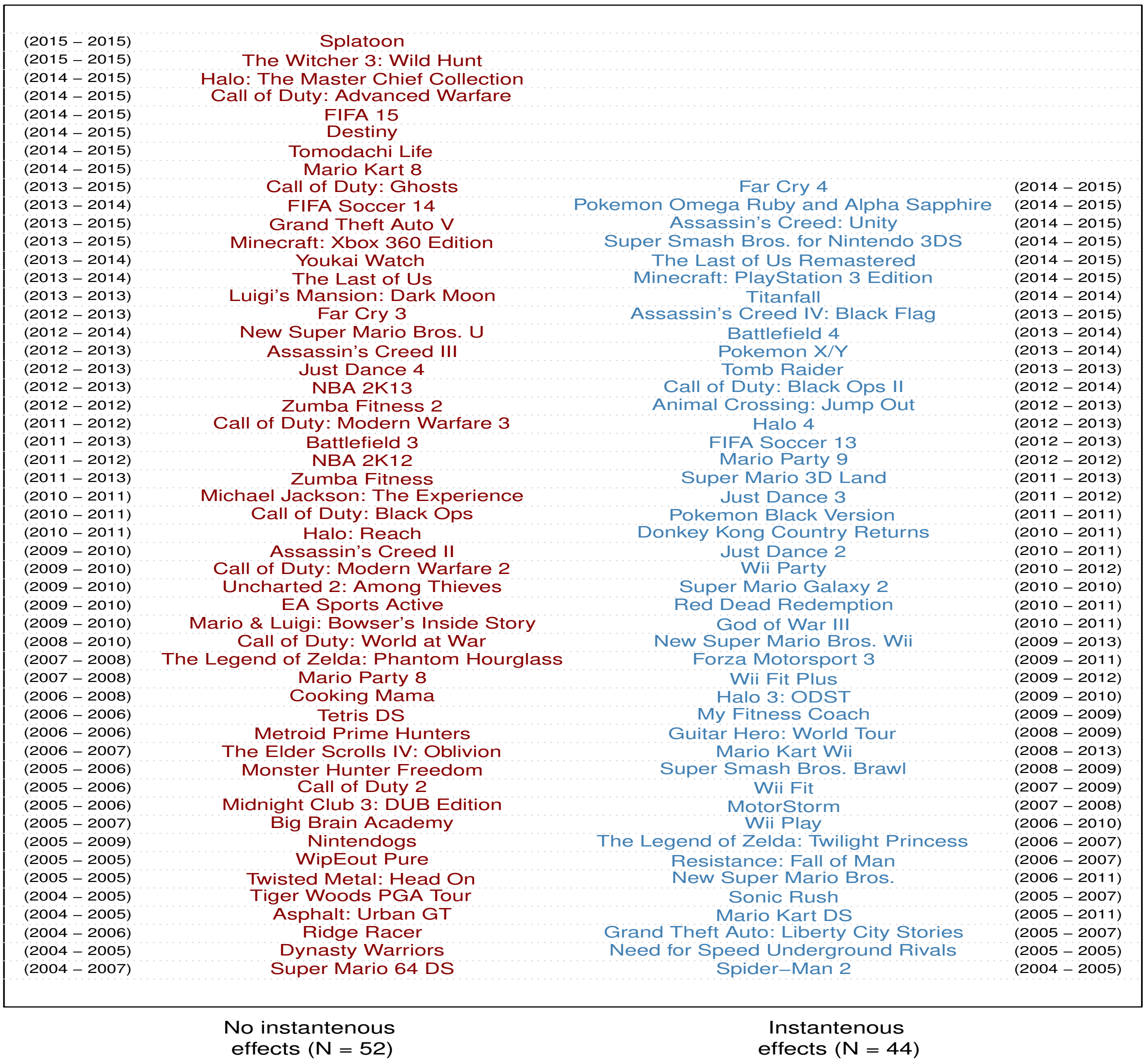

Fig. 10 Breakdown by Instantaneous Effects (observed calendar time periods in parenthesis on left and right of the groupings)

els improves only modestly from the inclusion of search volumes – and for many games, the inclusion actually reduces the performance. In other words, the results justify also a healthy dose of skepticism regarding the usefulness of search volumes in the context of video game sales modeling.

Also data limitations (see Section 6.3) explain a part of this divergence. Although a detailed examination was deliberately restricted outside to the scope of this paper, these cases may relate also to different cross-sectional abstractions. In particular, the results hint that the games that sell well may show stronger correlations between sales and search volumes. In this regard, it might be interesting to further use search data for examining the questions related to "*superstar*" games (Brynjolfsson et al 2010; Hyrynsalmi et al 2016; and for related EWOM and industry concentration impacts, see also King et al 2014; Williams 2002). There are also numerous other theoretically relevant abstractions (for a summary, see Marchand 2015, Fig. 3), including sequels and franchises (Situmeang et al 2014), popularity and life cycle stages (Zhu and Zhang 2010), platforms and genres (Marchand and Henning-Thurau 2013), and regional market differences (Chintagunta et al 2010). Although the results presented do not seem to vary across these additional abstractions, which is generally in line with some existing EWOM observations (Zhu and Zhang 2010), further research is required for thorough statistical assessments. The cross-sectional abstractions also lead to more complex correlation structures (Chintagunta et al 2010), although the results presented indicate that the longitudinal dependence structures

Table 6 Summary of Empirical Results

C_i	Support	Summary
C_1	Yes	Internet search volumes and video game sales volumes tend to exhibit only relatively short autoregressive "memories" that amount to less than two months.
C_2	Yes	Platform-specific, geographically varying, and often consecutive release timings are associated with high sales volumes.
C_3	Yes	Likewise, Internet search volumes are usually higher for the release weeks.
C_4	No	For the clear majority of observed games, past Internet search volumes do not notably affect the current video game sales.
C_5	No	For an even larger majority of the observed video games, sales volumes do not Granger-cause search volumes.
C_6	Partial	For many games, there is an instantaneous effect between the two volumes.

may still remain relatively simple. In other words, the presence of instantaneous effects is a good baseline also for further time series cross-sectional modeling.

The remaining key results are easier to summarize. The observed search and sales volumes exhibit only a relatively short memory (C_1). Both volumes are also visibly associated with the video game release timings (C_2 and C_3). Although no attempts were made to compare pre- and post-release periods, the results are generally in accordance with previous observations, particularly regarding the effect of product launches on publicity (Burmester et al 2015). Taken together, these empirical results provide weak support for the feasibility of using Internet search volumes for predicting video game sales with conventional time series techniques (the facet of methods). The results provide also some material for the practical evaluation facet.

6.2 Managerial Implications

The foremost managerial implication mirrors the foremost scholarly implication: the use of Google search volumes for foresight purposes is not a panacea.

When sales and search volumes do correlate strongly, the results indicate that the week-to-week relationship tends to be rather instantaneous; when sales peak, also searches tend to peak. Thus, for many games, the mass-scale Internet-wide electronic word-of-mouth advertisement reacts very quickly to changes in sales volumes, or the other way around. Consequently, it is reasonable to recommend that no long-lasting EWOM effects should be expected when marketing and advertisement strategies are planned for video games. Planning should also keep in mind that the scheduling of new releases for different platforms and geographical markets will impact not only sales volumes but also the amount of searches. As the results are relatively robust in this regard, the effect or release timings deserves a particular emphasis: it may be possible to optimize Internet-wide publicity with contextual release strategies. Although the results are less clear about the impact of such publicity upon sales *per se*, strategic planning may overcome the short life cycles and rapid diffusion among consumers.

With respect to diffusion itself, the empirical results support many of well-known scholarly observations about sales dynamics in the video game industry. Initial performance is crucial for later sales. Product life cycles are as important as these have always been.

When trying to infer about the short-run future video game sales, search volumes provide a valuable heuristic but with only limited statistical forecasting power, however. Rather than using search volumes for predicting video game sales, the results presented can be said to bespeak for using Google searches as a general "*quantitative trendspotting*" technique (Du and Kamakura 2012). In this respect, search volumes provide a valuable information asset for many practical foresight tasks, regardless whether search volumes improve the prediction of domain-specific time series. Given the general advice for collecting as much data about video game consumers as is possible (Chen et al 2016), it seems fair to summarize that also Internet search data should be incorporated into the data collection frameworks of game developers, publishers, and distributors.

Yet, it should be kept in mind that there are a number of limitations that hinder the use of Google search data, ranging from data quality issues to ambiguities in specifying proper search keywords. While big data analysis is possible with proprietary datasets obtained from Google's databases (Ginsberg et al 2009), the practical usefulness of the trend service remains limited also for business intelligence companies. For companies such as VGChartz, Ltd., the prospects of continuous real-time monitoring (cf. Carneiro and Mylonakis 2009; Dimpfl and Jank 2015) are hindered already by the lack of programmable data retrieval interfaces that could be easily incorporated into existing software solutions. But as data availability may improve in the future, an open mind should be kept for future possibilities, as always.

6.3 Limitations

The facet of data is best reserved for a number of important research limitations. To disseminate these data limitations, a few key points should be mentioned about both sales and search volumes. The points have also important research implications.

6.3.1 Sales

The notable theoretical limitations are related to different confounding factors, which, in turn, are imposed by different data problems. Besides the mentioned cross-sectional abstractions, there are also many longitudinal factors that may interfere with the results reported. For instance, changes in pricing are likely to affect sales volumes (Chintagunta et al 2009; Hyrynsalmi et al 2015), but such changes might affect also search volumes. The same applies to traditional, non-EWOM, marketing and advertisement *per se*. Ideally, these factors could be controlled in time series analysis, but, in practice, the availability of robust data imposes severe limitations upon the potential number of confounding factors that can be controlled in a single empirical study of video game sales (cf. Burmester et al 2015; Chintagunta et al 2009). Analogous argument applies for other limitations imposed by constraints for data availability and quality.

While weekly sampling frequency has been the "state of the art" in Internet search volume modeling (Drake et al 2012), it may be that the purchasing dynamics are more fine-grained particularly in the video game industry. The intra-week dynamics have likely also intensified in recent years, owing to the increasingly important electronic purchasing channels (Jöckel et al 2008), subscription services such as PlayStation Plus, game preorders, sequels, expansion packs, bonus deals, product bundles, and related contemporary video game industry characteristics. Many of these characteristics can be reasonably assumed to affect also search volumes, which likely vary not only in terms of weekdays but also according to working hours, for instance. Alas, currently, daily or hourly time series are available neither for S_t nor for G_t. While accepting these (external) limitations as almost universal constraints in the whole research domain, it is thus more relevant to focus on the specific (internal) data limitations specific to this study.

A noteworthy internal concern relates to the per-game sample sizes: due to data limitations, only rather short time series were available for many of the observed games. This concern can be accompanied by noting the interpolation solution (see Section 3.2.1), which is likely to cause some inaccuracies. This said, it is also worth noting that the use of data from VGChartz, Ltd. (2016) is seldom accompanied with acknowledgements about missing values and other related data quality issues (Marchand 2015; Situmeang et al 2014), and even when weekly sampling frequency is used (Kim et al 2014). In other words, this paper is far from being the only one affected by the data availability issues. For further scholarly research, a worthwhile goal would be to systematically survey and catalog different data limitations, which would also enable better understanding of the practical research potential in the domain. The surveying viewpoint does not need to be scholarly, however. When compared to some traditional industries, it may well be that the data limitations reflect much bigger practical issues, including the potential immaturity of business intelligence related to video games.

6.3.2 Searches

It is important to understand the limitations related to the empirical time series data provided by Google. Four significant limitations can be enumerated: (a) actual search volumes are not observable due to the indexing solution; (b) output is highly sensitive to the specific search terms used; (c) the trends vary according to the queried longitudinal period; and (d) geographical search parameters affect the output (Polgreen et al 2008; Reilly et al 2012). The first limitation affects directly some of the time series estimates. In particular, there is an undisclosed threshold that determines whether a search query is registered in Google's trend service. Consequently, a few of the time series estimates are affected by scenarios such as the one illustrated in Fig. 11. For this particular outlying video game, there is a long period of zeros in the indexed search volume, which does not mean that there would not have been searches for the game. While interpretation is one thing, the long period of zeros contributes also to the lack of strong instantaneous relationship between the series.

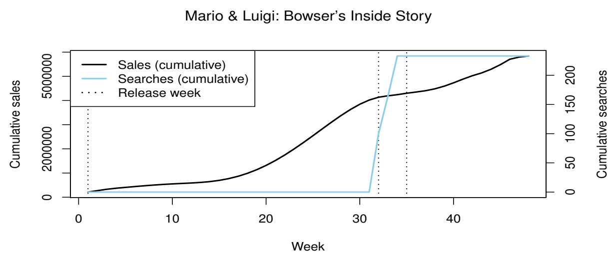

Fig. 11 An Illustration of Data Limitations

Not all of the four problems are equally significant as research limitations in this paper, however. While the limitations (c) and (d) were to some extent "controlled" by specifying a global scope and the same longitudinal period for all games, both problems are arguably external research constraints rather than internal limitations of this paper. Namely, the search volumes are insufficient for many games in some regions (see also Section 3.2.2), while Google's search data is available only up to 2004. Both issues limit the general applicability of Google trends for time series modeling of video game sales volumes. In particular, VGChartz, Ltd. (2016) provides data for sales in the United States, Japan, and Europe, but this geographic decomposition cannot be fully utilized due to the limited geographic coverage of Google trends for some games. Finally, the business and politics of search engines themselves should be also kept in mind (Jiang 2014). In other words, Google trends are presumably rather biased for observing Internet search volumes in countries such as China and Russia.

The same applies to the limitation (b), the general problems associated with specifying search keywords for research purposes. While it is rather difficult to specify suitable multi-language keywords for, say, fishing (cf. Wilde and Pope 2013), video game titles supposedly provide rather accurate queries from the EWOM perspective, especially when further enclosing the Google queries with quotation marks. Also game sequels and franchises can be controlled to a degree; while games such as Skate, DiRT, and Rock Band yield too general search terms, the sequels Skate 2 and 3, DiRT 2 and 3, and Rock Band 2 provide arguably sufficiently sharp search terms.

Analogous reasoning has been used for preferring abbreviations over company names (Drake et al 2012). Nevertheless, also the adopted strategy requires subjective exclusion of cases due to the prevalence of broadly entitled video games. Although reliability could be improved in this regard (for instance by evaluating the cross-agreement across multiple researchers, or crowdsourcing the searches; see Brynjolfsson et al 2016), it is generally difficult – if not impossible – to speculate about people's typical Google search term patterns without being affiliated with Google, Inc.[5] All in all, it thus seems safe to summarize that this essentially external "search term limitation" will continue to constrain scholarly Internet search volume modeling in general.

It is also possible to question whether the Google trends truly reflect the intended Internet-wide EWOM. For answering to this concern, it is first necessary to point out the general empirical landscape in the EWOM research domain. Arguably, most of the empirical literature derives hypotheses according to the availability of empirical data, which is then reflected against the endorsed branch of theoretical literature. Consequently, depending on the branch, the exactly same review data from a popular aggregation site (Metacritic), for instance, translates into EWOM effects (Hennig-Thurau et al 2015), software quality of games (Kim et al 2014), and even innovativeness (Storz et al 2015), resulting incoherence even regarding what is being observed.

Against this general research background, the use of search volumes arguably provides a better and more coherent way for observing mass-scale electronic word-of-mouth in the contemporary Internet. It is another question how useful this kind of mass-scale EWOM is for observing customer behavior. It may be that bigger is not always better in the video game industry. That is to say, the continuing importance of electronic purchasing channels leads also to richer and more robust databases. In this respect, companies such as Valve Corporation may already have a head start in practical foresight, given the popular Steam distribution platform for PC video games. Due to the described limitations associated with search volumes, it seems reasonable to argue that also scholarly research should focus on more nuanced data sources.

For studying diffusion among consumers, whether via purchases or EWOM, the existing aggregation sites, social media platforms such as Twitter or YouTube, and related sources seem more prolific than aggregated and indexed time series on Internet searches. Many of these sources would further enable the use of social network analysis for studying the actual human behavior underneath the EWOM concept. When compared to conducting a brief Internet search or glancing through a game review, which are both still relatively passive means for retrieving information, an active discussion on a social media platform may be more important for a purchasing decision. In addition to framing the focus closer to the original WOM concept (cf. Arndt 1967), studying such social network dynamics may deliver better understanding on how the associated network effects (Katz and Shapiro 1985) spread in the video game industry.

6.4 Concluding Remarks

The use of Google trends for applied time series modeling has sometimes been celebrated as a turn of tide for many practical research problems. This paper welcomes a healthy amount of skepticism for the commemorate insofar video game sales modeling is concerned. While Google trends offer a valuable instrument for observing Internet-wide EWOM effects, it is likely that search volumes provide only incremental forecasting improvements when benchmarked against the well-established diffusion models, for instance. This argument opens also a number of relevant questions for further research.

[5] It can be further remarked that some studies have tried to balance these issues (including the indexing problem) by scaling the observed search volumes with a generic high-frequency search term (Curme et al 2014). As said, however, it arguably remains unclear how useful such scaling is without knowing technical details about the proprietary Google search engine.

Besides (a) benchmarking of different time series models, (b) it would be relevant to evaluate whether the results are specific to the video game industry. Indeed, it may well be that the distinct characteristics of the video game industry are reflected in the results reported. The examples range from the so-called big bang dynamics and buzz generation tactics to the distinct network effects that characterize the video game industry. Although the results are not entirely contradictory with those obtained from the car industry (Hu et al 2014), for instance, it seems that the usefulness of Internet search volumes for sales modeling may generally vary from an industry to another.

Moreover, (c) it should be further evaluated whether the examined conjectures vary within the video game industry; it may be that the predictive power of search volumes increases through the potential combinatory effect with other variables. In particular, (d) further research is required for understanding the combined effect of Internet search volumes and the more traditional EWOM variables, including game reviews and the associated quantifications. In this respect, dimension reduction techniques, such as principal component analysis, may offer a valuable evaluation technique. Such techniques could also shed more light upon a question of (e) whether it is possibly to empirically remedy some of the elaborated data limitations associated with the publicly available Google trend data. Finally, (f) it would be worthwhile to carry out a more practical review regarding the potential usefulness of Internet search data for different software solutions. To give one example: the use of Google maps has become almost an ubiquitous trait of the contemporary world wide web, but the incorporation of search data is only seldom seen in third-party sites. If the current search engine frontier takes a turn towards open data in the future, also the practical prospects enlarge substantially.

References

Arndt J (1967) Role of Product-Related Conversations in the Diffusion of a New Product. Journal of Marketing Research 4(3):291–295

Athanasopoulos G, Guillén OTd, Issler JV, Vahid F (2011) Model Selection, Estimation and Forecasting in VAR Models with Short-Run and Long-Run Restrictions. Journal of Econometrics 164(1):116–129

Bårdsen G, Eitrheim Ø, Jansen ES, Nymoen R (2005) The Econometrics of Macroeconomic Modeling. Oxford University Press, Oxford

Brynjolfsson E, Hu Y, Smith MD (2010) Long Tails vs. Superstars: The Effect of Information Technology on Product Variety and Sales Concentration Patterns. Information Systems Research 21(4):736–747

Brynjolfsson E, Geva T, Reichman S (2016) Crowd-Squared: Amplifying the Predictive Power of Search Trend Data. MIS Quarterly (Forthcoming, published online)

Burmester AB, Becker JU, van Heerde HJ, Clement M (2015) The Impact of Pre- and Post-Launch Publicity and Advertising on New Product Sales. International Journal of Research in Marketing 32(4):408–417

Carneiro HA, Mylonakis E (2009) Google Trends: A Web-Based Tool for Real-Time Surveillance of Disease Outbreaks. Clinical Infectious Diseases 49(10):1557–1564

Chen G, Cox JH, Uluagac AS, Copeland JA (2016) In-Depth Survey of Digital Advertisement Technologies. IEEE Communications Surveys & Tutorials 18(3):2124–2148

Chintagunta PK, Nair HS, Sukumar R (2009) Measuring Marketing-Mix Effects in the 32/64 Bit Vide-Game Console Market. Journal of Applied Econometrics 24(3):421–445

Chintagunta PK, Gopinath S, Venkataraman S (2010) The Effects of Online User Reviews on Movie Box Office Performance: Accounting for Sequential Rollout and Aggregation Across Local Markets. Marketing Science 29(5):944–957

Choi H, Varian H (2012) Predicting the Present with Google Trends. The Economic Record 88(Issue Supplement s1):2–9

Cortada JW (2013) How New Technologies Spread: Lessons from Computing Technologies. Technology and Culture 54(2):229–261

Cox WE (1967) Product Life Cycles as Marketing Models. The Journal of Business 40(4):375–384

Curme C, Preis T, Stanley HE, Moat HS (2014) Quantifying the Semantics of Search Behavior Before Stock Market Moves. Proceedings of the National Academy of Sciences of the United States of America 111(32):11,600–11,605

De Nooy W, Mrvar A, Batagelj V (2011) Exploratory Social Network Analysis with Pajek, 2nd edn. Cambridge University Press, Cambridge

Dellarocas C, Zhang X, Awad NF (2009) Exploring the Value of Online Product Reviews in Forecasting Sales: The Case of Motion Pictures. Journal of Interactive Marketing 21(4):23–45

Dimpfl T, Jank S (2015) Can Internet Search Queries Help to Predict Stock Market Volatility? European Financial Management 22(2):171–192

Drake MS, Roulstone DT, Thornock JR (2012) Investor Information Demand: Evidence from Google Searches Around Earnings Announcements. Journal of Accounting Research 50(4):1001–1040

Du RY, Kamakura WA (2012) Quantitative Trendspotting. Journal of Marketing Research 49(4):514–536

Du RY, Hu Y, Damangir S (2015) Leveraging Trends in Online Searches for Product Features in Market Response Modeling. Journal of Marketing 79(1):29–43

Ginsberg J, Mohebbi MH, Patel RS, Brammer L, Smolinski MS, Brilliant L (2009) Detecting Influenza Epidemics Using Search Engine Query Data. Nature 457(19):1012–1015

Gluskin RT, Johansson MA, Santillana M, Brownstein JS (2014) Evaluation of Internet-Based Dengue Query Data: Google Dengue Trends. PLOS ONE 8(2):e2713

Google, Inc (2016) Google Trends, available online in January 2016: `https://www.google.com/trends/`

Granger CWJ (1988) Some Recent Developments in a Concept of Causality. Journal of Econometrics 39(1–2):199–211

Harada N (2007) Video Game Demand in Japan: A Household Data Analysis. Applied Economics 39(13):1705–1710

Hastie T, Tibshirani R, Friedman J (2011) The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer-Verlag, New York

Hendry DF, Juselius K (2001) Explaining Cointegration Analysis: Part II. The Energy Journal 22(1):75–120

Hennig-Thurau T, Wiertz C, Feldhaus F (2015) Does Twitter Matter? The Impact of Microblogging Word of Mouth on Consumers' Adoption of New Movies. Journal of the Academy of Marketing Science 43(3):375–394

Hewing M (2011) A Theoretical and Empirical Comparison of Innovation Diffusion Models Applying Data from the Software Industry. Journal of the Knowledge Economy 3(2):125–141

Hu Y, Du RY, Damangir S (2014) Decomposing the Impact of Advertising: Augmenting Sales with Online Search Data. Journal of Marketing Research 51(3):300–319

Hyrynsalmi S, Seppänen M, Aarikka-Stenroos L, Suominen A, Järveläinen J, Harkke V (2015) Busting Myths of Electronic Word of Mouth: The Relationship between Customer Ratings and the Sales of Mobile Applications. Journal of Technical and Applied Electronic Commerce Research 10(2):1–18

Hyrynsalmi S, Suominen A, Mäntymäki M (2016) The Influence of Developer Multi-Homing on Competition Between Software Ecosystems. The Journal of Systems and Software 11:119–127

Jiang M (2014) The Business and Politics of Search Engines: A Comparative Study of Baidu and Google's Search Results of Internet Events in China. New Media & Society 16(2):212–233

Jöckel S, Will A, Schwarzer F (2008) Participatory Media Culture and Digital Online Distribution – Reconfiguring the Value Chain in the Computer Game Industry. International Journal on Media Management 10(3):102–111

Juselius K (2006) The Cointegrated VAR Model. Methodology and Applications. Oxford University Press, Oxford

Katz ML, Shapiro C (1985) Network Externalities, Competition, and Compatibility. The American Economic Review 75(3):424–440

Kim J, Prince J, Qiu C (2014) Indirect Network Effects and the Quality Dimension: A Look at the Gaming Industry. International Journal of Industrial Organization 37(37):99–108

King RA, Racherla P, Bush VD (2014) What We Know and Don't Know About Online Word-of-Mouth: A Review and Synthesis of the Literature. Journal of Interactive Marketing 28(3):167–183

Landsman V, Stremersch S (2011) Multihoming in Two-Sided Markets: An Empirical Inquiry in the Video Game Console Industry. Journal of Marketing 75(6):39–54

Lütkepohl H (2005) New Introduction to Multiple Time Series Analysis. Springer-Verlag, Berlin

Lütkepohl H, Reimers HE (1992) Granger-Causality in Cointegrated VAR Processes: The Case of the Term Structure. Economics Letters 40(3):263–268

Marchand A (2015) The Power of an Installed Base to Combat Lifecycle Decline: The Case of Video Games. International Journal of Research in Marketing Published online in 15 July 2015

Marchand A, Henning-Thurau T (2013) Value Creation in the Video Game Industry: Industry Economics, Consumer Benefits, and Research Opportunities. Journal of Interactive Marketing 27(3):141–157

Pfaff B (2008) VAR, SVAR, and SVEC Models: Implementation Within R Package *vars*. Journal of Statistical Software 27(4):1–32

Piirainen KA, Gonzalez RA, Bragge J (2012) A Systemic Evaluation Framework for Futures Research. Futures 44(5):464–474

Polgreen PM, Chen Y, Pennock DM, Nelson FD (2008) Using Internet Seaches for Influenza Surveillance. Clinical Infectious Diseases 47(11):1443–1448

Polli R, Cook V (1969) Validity of the Product Life Cycle. The Journal of Business 42(4):385–400

Proulx R, Massicotte P, Pépino M (2013) Googling Trends in Conservation Biology. Conservation Biology 28(1):44–51

Rangaswamy A, Giles CL, Seres S (2009) A Strategic Perspective on Search Engines: Thought Candies for Practitioners and Researchers. Journal of Interactive Marketing 23(1):49–60

Reilly S, Richey S, Taylor JB (2012) Using Google Search Data for State Politics Research: An Empirical Validity Test Using Roll-Off Data. State Politics & Policy Quarterly 12(2):146–159

Ruohonen J, Hyrynsalmi S, Leppänen V (2015a) The Sigmoidal Growth of Operating System Security Vulnerabilities: An Empirical Revisit. Computers & Security 55:1–20

Ruohonen J, Hyrynsalmi S, Leppänen V (2015b) Time Series Trends in Software Evolution. Journal of Software: Evolution and Process 27(2):990–1015

Segev E, Baram-Tsabari A (2010) Seeking Information Online: Data Mining Google to Better Understand the Roles of the Media and the Education System. Public Understanding of Science 21(7):813–829

Situmeang FB, Leenders MA, Wijnberg NM (2014) History Matters: The Impact of Reviews and Sales of Earlier Versions of a Product on Consumer and Expert Reviews of New Editions. European Management Journal 31(1):73–83

Storz C, Riboldazzi F, John M (2015) Mobility and Innovation: A Cross-Country Comparison in the Video Games Industry. Research Policy 44(1):121–137

Su M, Rao VR (2010) New Product Preannouncement as a Signaling Strategy: An Audience-Specific Review and Analysis. Journal of Product Innovation Management 27(5):658–672

Suber-Jenkins A (2016) 'Pokémon Go' Is Losing Players Fast – Here Are the 4 Biggest Reasons Why, Tech.Mic, Available online in August 2016: `https://mic.com/articles/152459/pok-mon-go-is-losing-players-fast-here-are-the-4-biggest-reasons-why#.QQb65cr9X`

Vaughan L, Romero-Frías E (2014) Web Search Volume as a Predictor of Academic Fame: An Exploration of Google Trends. Journal of the Association for Information Science and Technology 65(4):707–720

VGChartz, Ltd (2016) Global Weekly Chart, available online in January 2016: `http://www.vgchartz.com/weekly/`

Vosen S, Schmidt T (2011) Forecasting Private Consumption: Survey-Based Indicators vs. Google Trends. Journal of Forecasting 30(6):565–578

Wilde GR, Pope KL (2013) Worldwide Trends in Fishing Interest Indicated by Internet Search Volume. Fisheries Management and Ecology 20(2–3):212–222

Williams D (2002) Structure and Competition in the U.S. Home Video Game Industry. International Journal on Media Management 4(1):41–54

Zhang Z, Zheng X, Zeng DD, Leischow SJ (2015) Information Seeking Regarding Tobacco and Lung Cancer: Effects of Seasonality. PLOS ONE 10(3):e0117,938

Zhou X, Li Q, Zhu Z, Zhao H, Tang H (2013) Monitoring Epidemic Alert Levels by Analyzing Internet Search

Volume. IEEE Transactions on Biomedical Engineering 60(2):446–452

Zhu F, Zhang XM (2010) Impact of Online Consumer Reviews on Sales: The Moderating Role of Product and Consumer Characteristics. Journal of Marketing 74(2):133–148