EXPLAINING BOX OFFICE PERFORMANCE FROM THE BOTTOM UP: DATA, THEORIES AND MODELS

by

Holly Russo
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
in Partial Fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Computational Social Science

Committee:

_______________________________________ Dr. Robert Axtell, Dissertation Director

_______________________________________ Dr. Richard Wagner, Committee Member

_______________________________________ Dr. Andrew Crooks, Committee Member

_______________________________________ Dr. Kevin Curtin, Acting Department Chair

_______________________________________ Dr. Donna M. Fox, Associate Dean, Office of Student Affairs & Special Programs, College of Science

_______________________________________ Dr. Peggy Agouris, Dean, College of Science

Date: _________________________________ Spring Semester 2016
George Mason University
Fairfax, VA

Explaining Box Office Performance From the Bottom Up: Data, Theories and Models

A Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

by

Holly Russo
Master of Arts
George Mason University, 2005

Director: Robert Axtell, Professor
Department of Computational Social Science

Spring Semester 2016
George Mason University
Fairfax, VA

COPYRIGHT

Copyright 2016 Holly A. Russo

All Rights Reserved

DEDICATION

This dissertation is dedicated to the late Dr. Charles K Rowley, Duncan Black Professor of Economics, President and Director of the Locke Institute. In the course of completing my B.Sc. in economics at GMU, I took several of Dr. Rowley's classes. Dr. Rowley encouraged me to apply to the Ph.D. program, something I had not previously even considered as a possibility. I regret that I did not finish before he passed, so he could see how much his recommendation influenced my path. But, I will never forget his influence, as well as his teaching and his tremendous kindness. If I think of the one person who truly changed the course of my career and life, it would be Dr. Rowley, and so I wholeheartedly dedicate this dissertation to him.

ACKNOWLEDGEMENTS

I am truly blessed to have such a wonderful support network around me, who saw me through this effort. I would like to thank Dr. Rob Axtell, my advisor, for never losing faith in me and talking me off the ledge when I lost faith in myself; Dr. Richard Wagner, for his enthusiasm and encouragement; Dr. Andrew Crooks, for offering much-needed support and meticulous editorial comments at the most crucial time. I would also like to thank Karen Underwood, who made me feel so welcome every time I came into the department, and who kept me from stressing about the process by being there to help. I owe a huge debt of gratitude and love to my amazing husband Russ, who woke me up with coffee every weekday at 0530 during the final year to support my work; who also listened to me lecture and ramble for long periods of time, being a sounding board and offering an outsider perspective, all of which helped me sort out some of the more thorny concepts and arguments. Also I want to acknowledge my beautiful step-daughters Sophia and Miranda with much love and best wishes for their bright futures, who both took an interest and offered their unwavering support. I want to thank Dr. Ron Morrison, who agreed to read, offer advice and coaching when I needed it, and who challenged me to improve my explanations. And I want to thank my Fairfax Harley Owners Group family of supporters and cheerleaders, who understood that I needed to step back from contributing to the chapter in order to focus on this work, and who always asked about my progress. No acknowledgement would be complete without the pets: my faithful cats Ding and Ling, who are 17 now and have physically occupied the seats next to me throughout all of my undergraduate and graduate work, benefitting from the heat of the laptop and snoring while I typed. And finally, I want to thank the army of guardian angels who kept me alive and out of trouble long enough to finish this, despite myself.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

2 stage least squares	2SLS
3 stage least squares	3SLS
Box office	BO
Coefficient of variation	CV
Long term box office	LTBO
Motion Picture Association of America	MPAA
Ordinary least squares	OLS
Rotten Tomatoes	RT
Standard deviation	SD
Short term box office	STBO
Word of mouth	WOM

ABSTRACT

EXPLAINING BOX OFFICE PERFORMANCE FROM THE BOTTOM UP: DATA, THEORIES AND MODELS

Holly Russo, Ph.D.

George Mason University, 2016

Dissertation Director: Dr. Robert Axtell

Every week, there are more than 50 movies playing in theaters from which movie-goers can choose. Analyses of the relative box office success of these films shows that it is Pareto-distributed, with roughly 20% of them earning 80% of the overall revenue. Arthur De Vany studied the potential causes of this 'winner-take-all' distribution through equation-based analyses, and theorized that the Pareto-distributed box office revenues we observe emerge from the micro-level complex adaptive behavior of movie-goers with imperfect information.

The literature is replete with attempts to explain how blockbuster movies occur, yet clearly none have found the secret formula. Several concluded a relationship between attributes and box office success through the information signaled or generated by the attribute, but all have suffered from an inability to separate the effect of movie attributes from the effect of word of mouth.

In this dissertation I created an agent-based model of movie attendance called ChatterMod, where movie-goers with heterogeneous and incomplete information choose from a supply of homogenous films devoid of 'preference' attributes (e.g. rating, genre, existence of stars) and then exchange information about the film with their neighbors. ChatterMod explains how the skewed macro-level distribution of box office revenues arises from a) a scale-free network of movie-goers with heterogeneous levels of movie awareness, b) each movie's revenue relative to existing films, c) the competition for movie-goer attention posed by information on new films, and d) the percolation of movie information throughout the network via word of mouth. In ChatterMod, there is no question as to whether or not a direct relationship between information and box office revenue exists, because there are no other attributes in the model that can affect that relationship. The model produced similar statistical behavior to prior empirical research and intuitive behavior when word of mouth and advertising were adjusted.

With this model, I provide a controlled laboratory for future research in movie information dynamics including the degree to which an initial movie-goer node's connectedness or each node's trust in its neighbors affect the success of a film, the identification of information tipping points or real world proxies for word of mouth, and the effect of the network structure on the overall distribution of revenue. This model can also be applied outside the film industry to study information dynamics in other areas such as music and book sales, website popularity, and voting. Finally, this dissertation provides further evidence that box office revenues cannot be predicted, and in the words of Goldman (1983) echoed by De Vany (2004), "nobody knows anything."

CHAPTER 1 - INTRODUCTION

What do movies and earthquakes have in common? Only 20% of each can be considered real 'movers and shakers', while the rest are mere tremblers (Bhattacharya & Chakrabarti, 2009; De Vany, 2004). For earthquakes this is a good thing; for movies, maybe not so good. We know we cannot control earthquakes, but filmmakers still hope that they can find the magic formula that will create films reaching a 10 on the Richter scale every time. Judging by the seemingly random nature of film success, they haven't found the winning formula yet.

Both the film blockbusters and earthquakes have been considered too random to predict (Goldman, 1983). However, statistical analyses have shown that frequency distributions of the magnitudes of earthquakes and box office success have both been found to be Pareto (power-law) distributed (Bhattacharya & Chakrabarti, 2009; De Vany, 2004; Sinha & Raghavendra, 2004) indicating that one or more definable processes are driving each phenomenon. One potential driver of box office revenues is the spread of information among a scale-free network of movie-goers, each with a different level of information about each currently-running film. The macro-level phenomenon we observe is the Pareto distribution of box office receipts, and the micro-level behavior producing that distribution is the passage of movie information from one node to the next in a network and movie-goer behavior in response to receiving that information.

The literature is replete with attempts to explain how blockbuster movies occur; . Many studies focused on the effect of movie attributes (Albert, 1998; Basuroy, Chatterjee, & Ravid, 2003; Moon et al. 2009; Neelamegham & Chintagunta, 1999 to name a few), and some concluded a relationship between attributes and box office success through the information signaled or generated by the attribute (for example Hennig-Thurau et al. 2007; Karniouchina, 2011), but neither those studies nor ones focused solely on information were able to collect a comprehensive real-world data set of this information. One study came close (Liu, 2006), by using the buzz on a movie-focused social network as a proxy, but there is no way to avoid the bias in such a dataset against word of mouth transmitted outside of electronic means. And with respect to buzz, there is not a consensus as to whether or not valance matters. Most importantly, all of these attempts have suffered from an inability to separate the effect of movie attributes from the effect of word of mouth.

Despite all of this research attention, there still has not been one combination of factors shown to comprise the secret formula for box office success. The proof is in the fact that we still have a heterogeneous supply of movies with a wide range of attributes from which only a few become major hits.

The model presented in this dissertation - ChatterMod - produces a complete data set of neutral information transmission about a complete population of films without attributes, and demonstrates how the skewed macro-level distribution of box office revenues arises from the micro-level transmission of information among a network of movie-goers, the timing of that transmission and the competitive environment in which it occurs. In ChatterMod, people share information with their neighbors about movies they have seen, and attend movies about which they have the most information. The model demonstrates how the factors a) each movie-goer's level of movie awareness, b) each movie's revenue relative to existing films, c) the competition for movie-goer attention posed by new films, and d) movie information percolating through a scale-free network via word of mouth combine to create an environment where only a few films are able to attract enough attention throughout their run to break away from the rest and become successful, resulting in the Pareto distribution of revenue we observe in real life. In ChatterMod, there is no question as to whether or not a direct relationship between information and box office revenue exists, because there are no other attributes in the model that can affect that relationship.

In past studies, models of movie industry economics have been empirical: generally econometric models examining the effect of several independent variables such as star-power, budget and genre on box office receipts (Elberse & Eliashberg, 2003). However, these models suffer from some limiting assumptions that are, in fact, far from the truth of movie industry economics and indeed many other complex adaptive systems. For one thing, complex adaptive systems are path-dependent resulting in auto-correlation in the error term, which is a violation of regression analysis assumptions. Another assumption important to most econometric analysis is that the data are Gaussian (normally) distributed; however box office receipts (De Vany, 2004; Sinha & Raghavendra, 2004) exhibit a Pareto distribution rendering econometric parameter estimates biased. Finally, economic analysis of supply and demand relies on the notion of system equilibrium; however the nature of a complex adaptive system is that it is constantly evolving with no equilibrium.

De Vany (2004) concluded that one could not predict movie success on star power, genre, budget or rating, and in fact any predictions made about movie revenue prior to release will have infinite errors, explaining that:

> *"Box-office revenue dynamics follow a Lévy stable process (which is a Bose–Einstein process with big leaps) and are asymptotically Pareto-distributed with infinite variance. This means that when a studio predicts a movie's revenue before it is released, the error is essentially infinite. A prediction that Movie A will gross X million plus or minus infinity, is no prediction at all."*

De Vany (2004) theorized that it was not the individual characteristics of the film itself that made the difference in success, but rather the complex process of the effect of those characteristics on the population and how that disseminated through the network.

Gemser et al. (2007) suggested that it is the volume of information overall, rather than the valance (positive vs. negative) information that drives ticket sales. It is not safe to say that only positive viewer feedback will prompt others to buy tickets, therefore we cannot necessarily say that subsequent ticket sales are higher because more previous viewers enjoyed the film, only that more previous viewers talked about the film. How many times have you found out about the existence of a film after it has closed in theaters, which you would have gone to see had you known about it? During the week of

March 18th 2016 there were roughly 122 films showing in one or more theaters, per the website Box Office Mojo (2016). How many can you name?

The casual observer can attest to the lack of significant relationships between movie attributes and box office revenues to some extent. Why did March of the Penguins, a documentary with Morgan Freeman narrating and no other actors at all, do so well? The well-known movie rating website Rotten Tomatoes™ (2015) gave it a Tomatometer™ score of 94% representing the percentage of positive critic reviews it received; it ranked #2 of all documentaries ever was the 27th highest grossing film of 2005. Conversely, why didn't Mars Attacks (my personal favorite) do nearly as well with an all-star cast including Jack Nicholson, Glenn Close, Pierce Brosnan, Michael J Fox, Annette Benning, Michael Short, Danny De Vito, Tom Jones, Jack Black, Christina Applegate, Sarah Jessica Parker, Rod Steiger and Natalie Portman? Rotten Tomatoes™ gave it a score of only 52%, and it only reached #39 for 1995, and #1,936 for all domestic films.

Those are two examples where the sentiment of the film – and the resulting box office performance – did not match the on-paper attributes of the film given what we would intuitively expect. Some research in this area implies that these are anomalies (see for example Albert 1998, Basuroy et al. 2003, and Hennig-Thurau et al. 2007) while others disagree (see for example Ravid 1999, and Liu 2006).

De Vany (2004) found that film revenue cannot be predicted by preference attributes or indicators such as casting, genre, rating, etc. and this evidence seems to support that assertion. Such little correlation between film opinion and performance begs the question: how are box office revenues influenced, if not by preference attributes? De Vany (2004), as well as Sinha and Raghavendra (2004) found that box office revenues followed a Pareto distribution where 20% of the films earn 80% of the revenue. This distribution renders attempts at predicting revenue for any single film - as well as the use of film attributes for attempting such a prediction - futile. De Vany theorized that it was the complex information dynamics among movie-goers that produced such a statistical distribution.

The thesis of this study is that similar macro-level Pareto distributions of box office revenues we see in real-life can be generated by a model that represents the micro-level interactions of movie-goers who choose from a supply of homogenous films based solely on the level of information they have about the existence of each film. By stripping out all of the 'preference' factors and modeling only an individual's awareness of a film through advertising and word of mouth, I show that box office performance depends on how much a person knows about a film rather than what a person knows about it, and the behavior that results from the spread of that information.

It is important to note that such a result does not conclusively prove that film revenues are based solely on awareness; it merely allows us to consider this as one possible explanation and explore it further. Intuitively, this explanation does make sense. Whether the information is advertising or word of mouth, regardless of valance it is only able to influence potential viewers if it is transmitted *and* received. Consider three potential cases: 1) Hollywood finally finds the perfect formula for a successful mega-blockbuster, but no one reviews it and it is not advertised at all; furthermore, no one talks about it at all after they've seen it. 2) The studios produce a mediocre movie that is heavily advertised and generates much buzz due to some feature such as special effects, or even negative buzz due to controversial themes. 3) A film is released for which there is a heavy amount of advertising and buzz, but the message is drowned out by some other noise that attracts attention away from that film's information. The first two are the extreme cases of no information vs. information overload: it is reasonable to assume that the great film with no information will suffer against the heavily promoted mediocre film. The third case is a little different: there is much information being transmitted but it is not being received.

It is important to note that ChatterMod has been developed to explore the relationship between information dynamics and box office revenue, not to predict what makes a particular movie successful or even what the total industry will achieve. My goal with ChatterMod was to grow the Pareto distribution of box office revenue from the ground up in order to describe the process by which this occurs.

So, how does a study that lends further evidence to the assertion that 'no one knows anything,' benefit Hollywood producers, as well as the economists and financial analysts trying to help them be successful? Perhaps it is to convince studios to rethink their strategy, focusing less on what they think movie-goers will like and more on how to get movie-goers talking. De Vany (2004) recommended that Hollywood adopt a portfolio approach to film production, rather than focusing on one film at a time. Such an approach would provide a better risk management strategy than counting on one particular film to produce expected results. Development of that approach is beyond the scope of this thesis but would be a logical next step, and could feasibly be tested with ChatterMod. As for the scope of this study, I discuss my work towards accomplishing the following steps in the next chapters:

- Chapter 2: Perform a review of the literature on previous studies of the film industry
- Chapter 3: Construct a computational agent-based model of a movie-going population and available films
- Chapter 4: Test that model against real-world MPAA statistics on movie attendance and film lifecycle, as well as the findings of De Vany and other researchers
- Chapter 5: Discuss the conclusions with respect to how information flow affects movie-goer behavior and ultimately the lifecycle of a film.

I conclude with a discussion of further research that can be performed with ChatterMod.

——••••——

[Martian Translator Device]: *All green of skin... 800 centuries ago, their bodily fluids include the birth of half-breeds. For the fundamental truth self-determination of the cosmos, for dark is the suede that mows like a harvest.*

[General Decker]: *What the hell does that mean?*

Quote from the movie *Mars Attacks*,

During the 'first contact' scene in Pahrump NV

CHAPTER 2 - LITERATURE REVIEW

The film industry has attracted substantial attention from researchers in an attempt to identify the factors that affect box office performance. In general, these factors fall into a few categories: film characteristics (stars, genre, rating), business strategy (budget, distribution and exhibition) and information (word of mouth, advertising, critical reviews), with most of the research efforts focusing on the film characteristics themselves. Elberse and Eliashberg produced an informative table in 2003 summarizing many of those research efforts, reproduced in Table 2-1 (Elberse & Eliashberg, 2003). This review will expand on some of these references and examine some additional literature from that period, as well as more recent studies. Section 2.5 has been devoted to the research of Arthur De Vany, as this dissertation significantly extends his work.

2.1 - Movie attributes: stars, genre, rating

Most of the past studies focus on the attributes of films, such as the presence of well-known stars, the genre of the film (e.g. horror vs. comedy), and the ratings (e.g. R, PG, G) assigned by the Motion Picture Association of America (MPAA). Data on these attributes are the easiest to collect and analyze, as they are objective measures and readily available from several sources including Variety.com, IMDB.com and The-Numbers.com. Despite this research attention, there still has not been one combination of factors shown to be the secret formula for box office success.

Table 2-1: Elberse and Eliashberg film success study review prior to 2003

Hypothesis	Evidence in literature
The higher a movie's number of screens in any given week, the higher its revenues in the same week	Strong evidence for a relationship with weekly revenues (Jones & Ritz, 1991, Sawhney and Eliashberg 1996)
The higher a movie's production budget, the higher its number of screens in the opening week	No evidence for budget, but strong evidence for revenues (Litman, 1982; Litman & Ahn, 1998; Litman & Kohl, 1989; Prag & Casavant, 1994; Wallace et al. 1993; Zufryden, 2000)
The higher a movie's star power, the higher its number of screens/revenues in opening week	Contradictory evidence: strong (Levin & Levin, 1997; Litman & Kohl, 1989; Neelamegham & Chintagunta, 1999; Sawhney & Eliashberg, 1996; Sochay, 1994; Wallace et al., 1993), limited to no evidence (Austin, 1989; De Vany & Walls, 1996; Litman, 1983; Litman & Ahn, 1998; Ravid, 1999)
The higher a movie's director power, the higher its opening screens/ revenues	Some evidence (Litman, 1982; Litman & Ahn, 1998; Litman & Kohl, 1989; Sochay, 1994)
The higher a movie's advertising expenditures, the higher its opening screens/revenues	Some evidence (Lehmann & Weinberg, 2000; Moul, 2001; Prag & Casavant, 1994; Zufryden, 2000)
The higher a movie's critical acclaim, the higher its opening screens/revenues	Contradictory: no (Eliashberg & Shugan, 1997), strong positive with cumulative revenues (Eliashberg & Shugan, 1997; Jedidi, Krider, & Weinberg, 1998; Litman, 1982; Litman & Ahn, 1998; Litman & Kohl, 1989; Prag & Casavant, 1994; Ravid, 1999; Sawhney & Eliashberg, 1996)
A movie distributed by one of the majors opens on a higher number of screens than an indie	No evidence for screens. Contradictory evidence for revenues: with opening week (Neelamegham & Chintagunta, 1999), with cumulative (Litman, 1983; Litman & Kohl, 1989), none with cumulative (Litman & Ahn, 1998; Sochay, 1994)
The more positive the word of mouth for a movie in a given week, the higher # screens/revenues in same week	No evidence for either (Neelamegham & Chintagunta, 1999)
The weaker a movie's competitive environment in any given week, the higher # screens/revenues in same week	Some negative for cumulative revenues (Litman & Ahn, 1998; Sochay, 1994) and weekly revenues (Jedidi et al., 1998; Zufryden, 2000)
The more a movie plays in a high season week, the higher its revenues	Some evidence: positive relationship with cumulative revenues (Litman, 1982; Litman & Kohl, 1989), contradictory for positive with weekly revenues (Ravid, 1999; Zufryden, 2000)

The bulk of the attribute focus has been on the presence / absence of stars, perhaps because of the financial implications of hiring an A-list actor and the level of controversy over their value. Table 2-2 shows a listing of sample movie incomes for well-known stars' roles in select films, the corresponding box office (BO) revenues, and the percent of BO allocated to each actor's income.

Table 2-2: List of representative actor incomes

Actor	Film	Year	Inc*	BO**	%
Sandra Bullock	Gravity	2013	70M	274.1M	26%
Robert Downey, Jr.	The Avengers	2012	50M	623.4M	8%
Johnny Depp	Pirates of the Caribbean: On Stranger Tides	2011	35M	241.1M	15%
Leonardo DiCaprio	Inception	2010	59M	292.6M	20%
Johnny Depp	Alice in Wonderland	2010	40M	334.1M	12%
Tom Cruise	War of the Worlds	2005	100M	234.3M	43%
Brad Pitt	Ocean's Eleven	2001	30M	183.4M	16%
Tom Cruise	Mission: Impossible II	2000	100M	215.4M	46%
Bruce Willis	The Sixth Sense	1999	100M	293.5M	34%
Tom Hanks	Saving Private Ryan	1998	40M	216.4M	18%
Mel Gibson	Lethal Weapon 4	1998	30M	130.4M	23%
Leonardo DiCaprio	Titanic	1997	40M	658.7M	6%
Tom Cruise	Mission: Impossible	1996	70M	181.0M	39%
Tom Hanks	Forrest Gump	1994	70M	330.3M	21%
Jack Nicholson	Batman	1989	60M	251.2M	24%

*** Source: ("List of highest paid film actors," 2014)**
**** Source: ("Box Office Mojo," 2014), US box office figures**

This table shows that the percentage of box office revenues taken by these actors fluctuates but is disproportionately high for some of them. For instance, Tom Cruise received 43% of box office revenue for *War of the Worlds* which only grossed $234M and 46% of revenues for *Mission: Impossible II* which grossed $215.4M , while Robert Downey Jr. only received 15% for *The Avengers* which grossed $623M.

Some authors indeed claim to have found empirical evidence that the presence of A-list actors has a significant effect on BO revenues. Albert (1998) found that the existence of well-known stars had a significantly positive effect on revenues, but as a signal of a particular film type informing potential viewers. Basuroy et al. (2003) also found that stars had a positive effect on revenue, but only for films that received negative critical reviews, and had little effect on films that received positive reviews. This evidence also points to the role of stars as a signal, in this case the movie-goer has two signals to choose from – the reviews and the presence of a star – and chooses the latter over the former as if he trusts the star more than the critic to convey accurate information about the film. Others who find significant positive statistical relationships between stars and box office revenues include Canterbury et al. (2001), Collins et al. (2002), Hennig-Thurau et al. (2007), and Neelamegham et al. (1999).

Several researchers have found no significant effect on box office revenues directly associated with particular movie attributes, but that the attributes had an indirect effect. Ravid (1999) identified that any significant increase in budget – either for stars or anything else – has a positive effect on revenues. Similarly, Elberse and Eliashberg (2003) found that stars and other movie characteristics influenced box office revenues indirectly through exhibitor screen allocations which they posited was a signal of exhibitor confidence in the film (thus quality) to the movie-goer. Karniouchina (2011) found that stars do indirectly have a positive impact on revenues from their ability to generate buzz and subsequently crowds for opening weekend, however Liu (2006) found this not to be the case.

Hennig-Thurau et al. (2007) actually found that A-list stars and directors exhibited a *negative* effect on box office revenues when considering both direct and indirect effects. They attribute this to either the tendency to compensate for a poor quality film with big names, or the higher expectations that movie-goers hold when those big names are involved.

Although opinions differ, the general consensus here is that the mere existence of a star in a film alone is not necessarily a recipe for box office success. With that in mind, I now turn to the business side of the film industry.

2.2 - Business strategy: budget, distribution and exhibition

The primary links in the film business chain are the producer / studio, the distributor and the exhibitor. Some studies of box office revenue have focused on film production budgets, however budget information that is accurately comparable across films is difficult to obtain due to things like profit participation by actors and directors (King, 2007). Even with a "well-trusted source used by industry decision makers," Cabral & Natividad (2014) were only able to obtain budget data for about 79% of their study sample and were required to apply industry estimation methods for the remainder. Despite this, they and others claim to have identified the effect of budgets on revenues and as with other film-specific factors, the findings differ.

Some found that budgets had no significant effect on movie revenues (Liu, 2006) while others such as Moon et al. (2009) observed that budgets actually had a negative effect on revenues after accounting for all other movie-related factors in their model. Conversely Ravid (1999) found while studying the effect of big name stars on the success of a film that any big investment in production – including a big name star – had a positive effect on box office revenues.

There is literature pointing to the absence or presence of other factors influencing the effect budgets have on revenues. Basuroy et al. (2003) found that budgets do have a positive effect on film performance, but only in the presence of negative critical reviews. His findings are supported by Pangarker & Smit (2013) who found that budget had a significant positive correlation with revenues (the most significant of their factors), and a significant *negative* correlation with critic reviews (expressed here on a scale of 0 to 100 with higher numbers being better reviews). The correlation with reviews implies that as review ratings rise the effect of budgets falls.

Others indicate that budgets may only be a part of the causal chain, affecting revenues through the other factors they influence. For instance, Karniouchina (2011) found that big advertising budgets were associated with increased movie buzz in Yahoo Movie data; both advertising and buzz contribute to the significant increase in opening week revenue that Muser (2011) observed as positively correlated with production budgets. Prior to both of these studies, Hennig-Thurau (2007) found that big budgets were positively correlated with revenues, but that greater advertising and wider opening 'resulted' from these budgets and that both of those contributed to revenues.

Once the film is produced, the distributor is responsible for determining the release pattern (e.g. wide, platform and limited release), the theaters and timing of release. The distributors base this strategy on their expectations of demand for the film, which are in turn based on their experience with prior films and knowledge of the industry environment. With this strategy in mind, the distributor contracts with one or more exhibitors to show the film at their theater for a specified period of time. Research into the relationship between this strategy and revenues focuses on the number of screens on which the film is exhibited within a geographical area.

This is the area where there seems to be the most consensus with respect to the significance of the effect on revenue. Neelamegham et al. (1999) found that number of screens on which a movie was released was the most significant predictor of success – in their estimate, "almost nine standard deviations greater than zero". Elberse and Eliashberg (2003) also found a significant correlation between the number of screens and box office performance, asserting that other movie attributes affect the number of screens, thus affecting performance indirectly. With Shugan, Eliashberg also found that the number of screens was an important predictor throughout the life of the film. (Eliashberg & Shugan, 1997). Both Basuroy et al. (2003) and Albert (1998) found that the number of screens had the most significant impact on revenues week to week, above factors such as star power, budget and reviews.

Note that in several of the studies on the effects of movie-specific attributes on revenue, authors have noted the possibility that the most direct influence on revenues is the information that results from budgets, stars, screens, etc. (Albert, 1998; Basuroy et al., 2003; Elberse & Eliashberg, 2003; Hennig-Thurau et al., 2007; Karniouchina, 2011; Pangarker & Smit, 2013; Ravid, 1999). Most of these focus on more indirect communication: the signals of quality that are sent to the potential movie-goer based on the existence of stars, the number of screens or the size of the budget. There is another area of literature, however, that focuses on more direct forms of communication: advertising, reviews and word of mouth (WOM). I now focus on this area of literature.

2.3 - Information: critical reviews, word of mouth, advertising

Until a movie-goer actually sees the film, she must rely on information from other sources to learn of the film's quality and characteristics. But even with this information, she will not be able to form her own valuation of the film prior to seeing it. Films are *experience goods*, a term attributed to Philip Nelson (1970) referring to goods for which the buyer cannot ascertain the true value prior to consumption. This falls between search goods, for which the buyer can determine the value, and post-consumption goods for which the buyer may still have trouble with a valuation even after consumption. Because of this, information such as reviews, advertising and WOM play a significant role in the potential viewer's decision process. There is a variety of studies that focus on movie information and the characteristics of that information: valance, volume and source.

Studies that focus on information valence and/or volume differ as to which has the greater influence and under what circumstances. For example, in studying the effect that film critics have on the success of a film an interesting discovery was made by Basuroy et al. (2003) who identified that valance had some significance in that films with more negative reviews than positive ones were helped by other factors such as star power and budget, but the same was not true where positive reviews outnumbered negative ones. They also found that the effect of positive reviews lingered while the effect of negative reviews declined over time. Basuroy's results differ from Eliashberg and Shugan (1997), who found that reviews were predictors (not influencers) of long term box office (LTBO), and did not affect or predict short term box office (STBO) at all. In their study STBO was affected more by trailers, advertising and number of opening screens than reviews.

In a study that Eliashberg performed with Sawhney (1996) to construct a revenue forecasting model called BOXMOD-I that included movie-goer behavior, he found that films which were heavily promoted through trailers and advertising resulted in exponentially declining box office revenues and very short average time-to-decide. Less promoted films that relied more on word of mouth showed longer time-to-decide and non-monotonic revenue patterns represented by the Erlang-2 or Generalized Gamma distributions.

Elberse and Eliashberg (2003) found that advertising was key for opening week performance while WOM was key for subsequent weeks. However, since the authors used prior week's performance as a proxy for WOM they actually demonstrated that current week's performance depends on prior week's performance, which logically would not hold for opening week. Therefore their study calls into question whether WOM or prior week performance is key, both of which have been purported to affect revenues in other literature. Neelamegham and Chintagunta (1999) tried the same approach to estimating WOM, using cumulative viewers as a proxy, but discarded the variable as it was not supported empirically. Regardless, the results of the study by Elberse and Eliashberg (2003) support the later research of Hennig-Thurau et al. (2007)

who also found that box office actions played a significant part in opening week revenues but ‘quality’ of film became more significant in the following weeks as more objective information about the film (e.g. word of mouth) began to spread.

The Hennig-Thurau (2007) study implies that valence of buzz is significant: if quality of a film is positively correlated with LTBO then one would assume this is because positive buzz results from quality, increasing potential viewer awareness and interest. Likewise, negative buzz has also been shown to be significant: Mizerski (1982) demonstrated the effect of buzz valance in a laboratory setting, where negative reviews from friends who had previewed a film affected the subsequent movie-goer’s opinion of the film. Burzynski and Baker (1977) observed the same thing in the field, where negative discussion of those leaving a film influenced those waiting in line to see the film to ask for their tickets to be refunded.

More recent studies on volume and valance of WOM have the advantage of WOM expressed on social media, which was not available to researchers conducting studies in the 80s and most of the 90s. One such study on WOM by Liu (2006) made use of WOM data collected from the Yahoo Movies web site. Although social media data is not a perfect sample, it is a much more comprehensive and representative sample than the lab data of prior studies. Liu found robust results demonstrating that volume of WOM both pre- and post-release was a statistically significant influence on BO, while the valance of WOM was not.

In a 2014 study, (Cabral & Natividad, 2014) the authors take a different approach and examine the effect on overall revenues of being number 1 at the box office The authors find that achieving the top of the box office is positively associated with a significant boost in total box office, and that the effect is greater for movies that were not heavily advertised prior to release. They assert that their findings provide evidence that being #1 raises potential movie-goers' awareness of the film.

In the past few sections I outlined the wide variety of factors found to be significant in box office revenues, and the circumstances under which they influence movie-goers, existing in the literature. I devote the next section to reviewing the assortment of methods used in analyzing the film industry.

2.4 - Past modeling approaches

Researchers focused on the film industry have employed a plethora of approaches including interviews, laboratory experiments and equation-based modeling.

Hierarchical Bayesian Poisson Regression: Neelamegham et al. (1999) developed a Bayesian modeling framework to predict weekly viewership, and estimated the parameters using Markov Chain Monte Carlo sampling. They chose a Bayesian model because it produced a predictive distribution rather than a point estimate of the box office performance of a film, as well as a measure of uncertainty. Their data set consisted of movie performance in 13 different countries. They used cumulative viewership as a proxy for word of mouth, as well as stars (presence or absence), genre, cumulative weeks, and a dummy variable for independent vs. major studio. In estimating the parameters of the Bayesian model, the modelers assumed that intercept and covariates were mutually independent, however in reality this can't be the case. Primarily this is because aside from contractual agreements with distributors, the length of time a film runs is based significantly on its cumulative viewership which in this model is used as a proxy for word of mouth, thus it is likely that these two variables are auto-correlated. They did decide that cumulative viewers wasn't a good proxy for word of mouth as the estimated parameter was negative, and instead concluded that it represented market saturation effects. In addition, there is likely a feedback loop between 'trend' (cumulative weeks) and revenue, as the revenue from previous weeks affects the number of weeks the film will run (per the decision of the film renter). Revenue affects whether or not the movie will run for another four weeks, and the run affects the potential revenue. This type of feedback loop can't be efficiently represented with equation based models.

Queuing Theory: Sawhney and Eliashberg (1996) used a queuing theory framework to conceptualize the movie adoption process (time to decide to see the movie and time to act by going to see the movie). They extended the theory by allowing for a delay between decision and action due to distribution limits and other factors. With the resulting model BOXMOD, they had some success predicting cumulative box office success (scale) using historical data, but much more difficulty predicting the temporal pattern of revenues (shape).

Markov Chains: As a follow on to their 1996 study, Sawhney and Eliashberg, with two others, implemented a Markov Chain model of consumer information and adoption behavior, called MOVIEMOD. (Eliashberg, Jonker, Sawhney, & Wierenga, 2000). In their model, consumers experienced states of undecided, considerer, rejecter, positive spreader, negative spreader and inactive. State transitions are triggered by advertising and word of mouth; probability of exposure to advertising depends on spending, and the probability of exposure to word of mouth depends on the number of spreaders, frequency of interaction and time-span of active spreading. The authors calibrated the model using lab interviews of movie-going and communication behavior. A diagram of their model is shown in Figure 2-1.

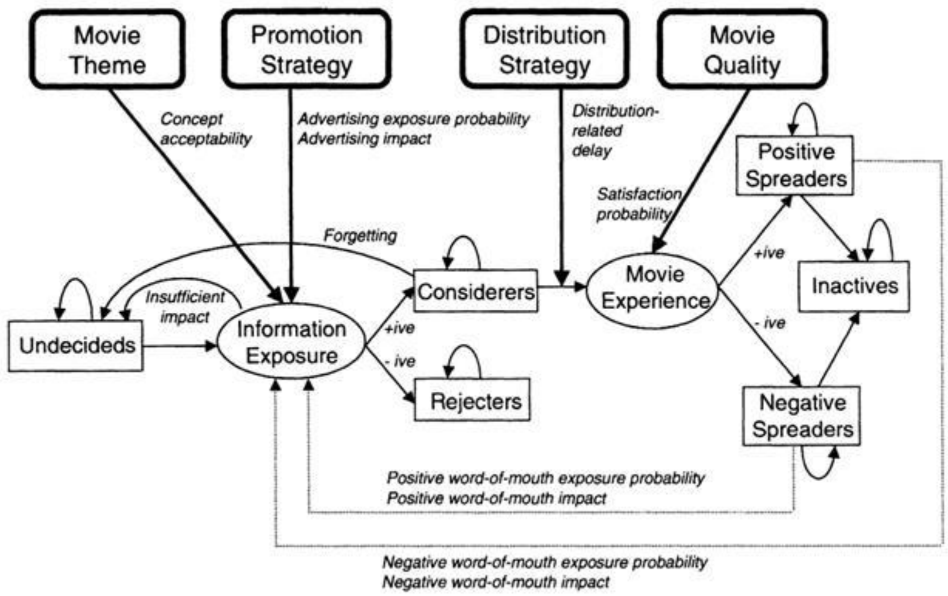

Figure 2-1: Representation of adoption process in MOVIEMOD. Source: Eliashberg et al. (2000). MOVIEMOD: An Implementable Decision-Support System for Prerelease Market Evaluation of Motion Pictures. Marketing Science, 19, 226–243

As you can see, MOVIEMOD does not account for any movie-specific factors other than theme; instead it focuses on information diffusion. However, MOVIEMOD does not incorporate the existence of multiple competing films, and was calibrated on only three films (Groundhog Day, Shadow Conspiracy and The Cemetery Club). Nonetheless it does provide a useful framework for information diffusion in ChatterMod, as will be discussed in Chapter 3.

Dynamic Simultaneous Equation Models: Continuing his work with various types of box office forecasting models, Eliashberg worked with Elberse (2003) to develop a structural equation model accounting for feedback loops and avoiding the issues with regression assumptions of independence. The important advance in this model was that the feedback loop relationship between revenues and number of screens was modeled, such that screen allocations affected revenue and prior week's revenue affected screen allocation decisions for the subsequent weeks. To estimate the model the authors applied ordinary least squares, 2-stage least squares and 3-stage least squares methods.

The data set the authors used consisted of 164 movies that were produced in the US in 1999 and appeared in the top 25 at least once. The data they obtained for those movies was weekly revenue and number of screens. They also incorporated word of mouth, however they used prior week revenues as proxy and asserted that this was in accordance with research by De Vany & Walls (1996), Hahn et al. (1994), Lilien et al. (1981), and Moul (2001). Unlike previous research, the authors also included variables for competitor movies, differentiating between new releases weighted by budget and top 25 movies weighted by age.

Probit and Logit Models: Several studies employed probit and/or logit models as a way to estimate parameters in the presence of the heavy-tailed distributions found in box office revenue data, as an alternative to OLS. Indeed, the study by Sinha and Raghavendra (2004) is just one of several studies that found a power law in rank distribution of gross revenues "for the most popular movies" with exponent close to -0.5.

Collins et al. (2001) performed standard econometric analysis on the UK movie industry and found that revenues exhibit an unbounded variance, which “undermines much of the existing work relating a film’s performance to its identifiable attributes within an OLS model.” The authors used a sample of 216 films that opened *and* closed between January and November 1998. They then apply probit and logit models (similar to De Vany and Walls, 1996), turning revenue data into binary data and estimating the likelihood that a movie will become a blockbuster (Collins et al., 2002).

Liran (2007) applied a nested logit demand model in order to estimate the underlying market size and its change with respect to seasons and/or supply. His sample of films consisted of all films released 1985-1999 (initially 3523 films) filtered down to films that at some point appear on 600 or more screens and other criteria, with the final sample being 1956 titles. Also, Liran considers the 'opening week' of his sample films to be the first week the film appeared on 600 screens, counting weeks prior to that as 'advertising'.

On a much smaller scale, Neelamegham and Jain (1999) studied the cognitive and emotional factors of consumer choice for three films: Little Big League, Forrest Gump and I Love Trouble. The authors employed a combination of lab experiments and a probit model study how things like consumer expectations and word of mouth affected choice and post-choice behavior. Their conceptual model is illustrated in Figure 2-2.

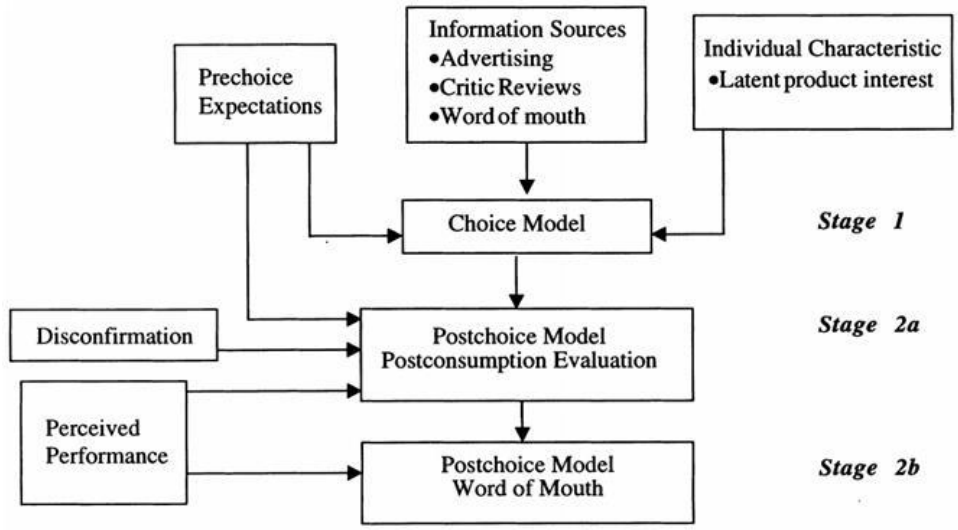

Figure 2-2: Consumer choice model overview. Source: Neelamegham, R., & Jain, D. (1999). Consumer Choice Process for Experience Goods: An Econometric Model and Analysis. Journal of Marketing Research, 36, 373–386

What is interesting in this model is the influence of the difference between pre-choice expectations and perceived performance on post-choice evaluation and behavior. So, it's not that the movie must be good, it just must meet the viewer's expectations to garner a positive evaluation and WOM referral.

Cellular Automata: Proykova and Stauffer (2002) constructed a social percolation model using cellular automata to represent heterogeneous agents with differing thresholds for quality, randomly distributed uniformly between 0 and 1. The model introduced one movie at a time, with quality based on that of the previous film: lower quality if the previous movie was a success and higher quality if it was a flop. The authors modeled diffusion of information from media advertising and word of mouth. Advertising only broadcasts until the first group attends the film, during which the ads inform some number of agents. Those with quality expectations below the advertised quality will go see the film, and inform others of what they saw.

This model assumes that all agents eventually learn of the film, because the process continues until all either go or decide not to go. In addition, the model does not account for competing films. Finally, the scale free network is a much more accurate representation of a population (Barabási & Albert, 1999) than the cellular automata paradigm. However, this model is notable in that it is closer to representing the complex adaptive processes driving box office revenue than the equation based models of other studies.

Structural Equation Modeling, Path Analysis and Latent Class Analysis: Hennig-Thurau et al. (2006, 2007) applied several causal analysis techniques including directed acyclic graphs for path modeling as shown in Figure 2-3, as well as structural equation modeling and latent class analysis on 331 films between 1999 and 2001.

Both of the studies by Hennig-Thurau et al. were designed to distinguish the factors that affect short term box office revenues vs. long term revenues, as well as identify both direct and indirect effects of revenue fluctuations. The issue particularly with the path analysis, is that it employs a directed acyclic graph which forbids the existence of feedback loops.

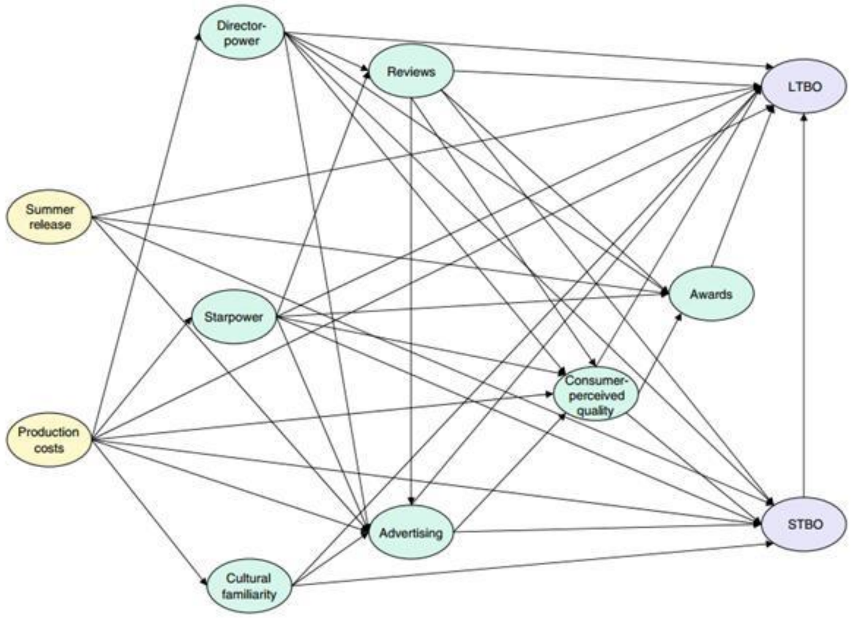

Figure 2-3: Model of relationships and influence on revenues. Source: Hennig-Thurau et al. (2007). Determinants of Motion Picture Box Office and Profitability: An Interrelationship Approach. Review of Managerial Science, 1, 65–92

Data: Just as important as the variety of modeling approaches used here, is that each of these studies involved sample sets of films that differed as well, along several dimensions. As described in the discussion above, some data sets included all films in the US (Sinha & Raghavendra, 2004), while others only included those in Variety's Top 50 (De Vany & Walls, 1996) and yet others only a small set of films (Eliashberg & Shugan, 1997; Liu, 2006). There were studies that included a sample of films from other countries (Basuroy et al., 2003; Ravid, 1999) while others focused only on the US (Sinha & Raghavendra, 2004). The number of years' worth of films varies from one to as many as 27 years (Cabral & Natividad, 2014; Liu, 2006), and because the studies cited here span several decades, the range of included production years does too from 1937 to 2010 (Dale, 1937; Pangarker & Smit, 2013).

Consider the possibility that the movie-going public has tastes that change year to year (or decade to decade); this isn't too far a stretch. For several years, the public might be more interested in seeing massive special effects in action films. Something changes and they begin to prefer more complex dramas with A-list stars. Then, something like Little Miss Sunshine comes along and turns the tide toward indie films that appeal to the whole family. This example demonstrates how sampling films from different decades can have different results with respect to what movie-specific factors – in this case budget, stars and MPAA rating respectively – have a more significant effect on box office revenues.

So, what have we learned from the review thus far with respect to the factors that influence box office revenues? De Vany (2004) summed it up well with the following, which has been quoted by several of the authors cited in this review so far:

> *"It would have been hard to imagine at the outset that by applying high-brow mathematical and statistical science we would end up proving Goldman's fundamental truth that, in the movies, 'nobody knows anything.'" (De Vany, 2004)*

2.5 - Arthur De Vany

In his quest to determine if it is at all possible to predict the success of a movie, De Vany studied the industry from many angles including revenues, profit, star power, genre and release strategies. His many published studies on the industry were collected and reproduced in his book titled **HOLLYWOOD ECONOMICS: HOW EXTREME UNCERTAINTY SHAPES THE FILM INDUSTRY** (De Vany, 2004). De Vany's conclusion was that multiple aspects of the industry were represented by the stable Paretian distribution with a finite mean driven by extreme events as well as an infinite variance driven by non-linear information dynamics among the movie-going population. It was this infinite variance that De Vany said supported Goldman's assertion that 'nobody knows anything', in that the errors on predictions would be infinite and measures of the precision of those predictions would calculate to zero.

Early in his work, De Vany sampled 300 films from Variety's Top-50 movies by week and identified that 20 percent of the films received 80 percent of the revenue. Plotting the rank-revenue relationship yielded the graph in Figure 2-4.

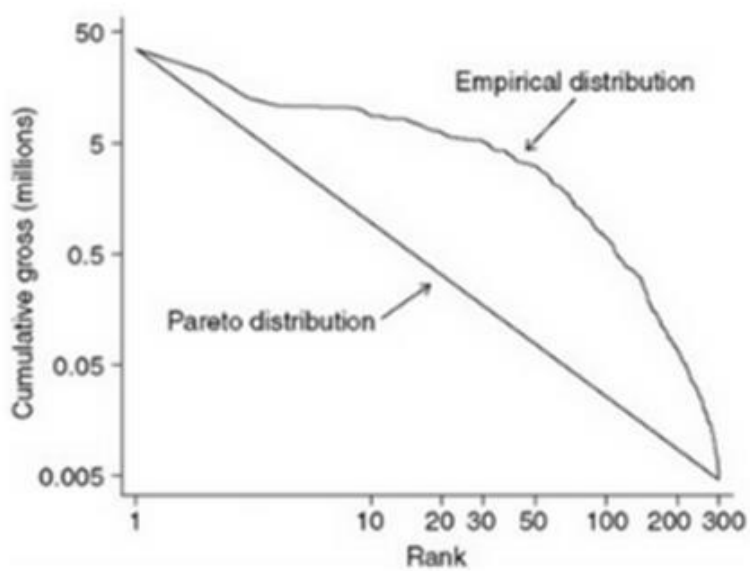

Figure 2-4: Rank-revenue relationship of films from Variety's Top-50. Source: De Vany, A. (2004). Hollywood Economics: How Extreme Uncertainty Shapes the Film Industry. London, England: Routledge

He applied statistical tests recommended by Nolan (1999) to reject the null hypothesis that the data were normally distributed and determine that the stable Paretian distribution was more appropriate. He also noted that the characteristic of self-similarity exists with an approximate tail weight of 1.5 across several aspects of the industry including:

- Cultural: in many countries including US, UK, Australia, Hong Kong and Ireland, a revenue/rank plot will exhibit the same distribution
- Temporal: revenue/rank plots over time exhibit the same distribution, as well as time series data of revenues
- Budget: separating movies into groups by budget (low, medium and high) results in all three groups exhibiting the same distribution by revenue/rank
- Also artist career length and pay

In observing this phenomenon, De Vany developed a statistical definition of 'legs' with respect to a film's run; to achieve 'legs' a film must gross enough revenue "to reach the upper Paretian tail of the distribution where the conditional expectation of forward revenue is linear in past revenue."

De Vany's subsequent studies focused on the adaptive behavior that was driving this phenomenon, and how to manage risk as a result of it. The main focus on drivers of this winner-take-all distribution of revenues in De Vany's studies was on the information dynamics of the movie-going population. In particular, he focused on the difference between non-informative vs. informative information cascades. A wide release (opening on many screens) can result in a larger non-informative cascade, where movie-goers are more likely to attend a film based on information they receive regarding massive opening weekend revenues. This initial cascade contains only 'quantity' data on the film, which De Vany showed to only be effective in the initial weeks of the film. He demonstrated a bifurcation of the revenue path of films after about five weeks as illustrated in Figure 2-5 such that some gain 'legs' and carry on to be successful while others begin to fail, and that gaining legs (thus movie-goer approval) was the most important factor in the monetary success of a film above other factors such as stars, genre and wide release.

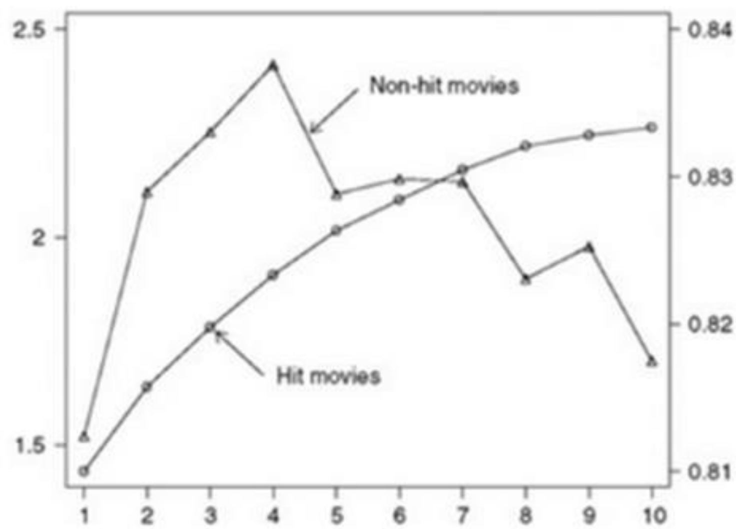

Figure 2-5: Bifurcation of revenue into hit and non-hit paths. Source: De Vany, A. (2004). Hollywood Economics: How Extreme Uncertainty Shapes the Film Industry. London, England: Routledge

This phenomenon he attributed to the informative information cascade, which contains more 'quality' data on the film, as the initial viewers begin to spread their opinions on what they've experienced via WOM. As a result, De Vany pointed out that a wide release is a risk in that it can propel a good film or kill a bad film as a result of the WOM effect of the informative cascade, rejecting the belief that a wide release is necessary or even sufficient for success at the box office. In fact, he observed that the opening box office data was only useful for predicting movies that would fail. De Vany's findings with respect to the significance of studio actions (such as wide release) in opening weeks followed by the significance of WOM in later weeks agree with those of Elberse & Eliashberg (2003) and Hennig-Thurau et al. (2007).

De Vany observed that film distributors and exhibitors can adapt *local* supply to updated *local* demand information by widening or narrowing the original release, shortening or lengthening the planned run, and even altering admission prices. This behavior represents the supply side of the constant feedback loop – or *complex adaptive process* – occurring between movie demanders (movie-goers) and suppliers (distributors and exhibitors) and driving movie revenues toward the stable Paretian distribution.

In light of his assertion that no one can predict the success of any individual film, De Vany recommends a portfolio approach to movie approval and financing. He states that although studios implicitly take this approach already by planning out their production activities into the future, they green-light films on an individual basis independent of other film decisions rather than analyzing and approving the portfolio as a whole. He concludes with the following which, based on his logic here, can be applied to any system that is complex and adaptive:

> *"Anyone who claims to forecast anything about a movie before it is released is a fraud or doesn't know what he is doing. The margin of error is infinite.*

That does not mean that he won't ever get it right, only that he seldom will and only because of sheer luck."

De Vany, 2004

——••••——

[Rumack]: *You'd better tell the Captain we've got to land as soon as we can. This woman has to be gotten to a hospital.*

[Elaine Dickinson]: *A hospital? What is it?*

[Rumack]: *It's a big building with patients, but that's not important right now.*

Quote from the movie *Airplane!*

CHAPTER 3 - MODEL DESCRIPTION

3.1 - Overview

Arthur De Vany (2004) used regression analysis to estimate many of the relationships in the film industry from a top-down perspective. This was an approach also adopted by many of the other authors cited in Chapter 2. For my study, I chose to build an agent-based model with a scale free network of movie-goers, to study the information dynamics from the bottom-up. Agent-based modeling is the most appropriate approach for this study because it can handle aspects of the movie industry information dynamics that other approaches cannot model, such as:

- *Time*: The timing of information transmission throughout the life of the film has been shown by several researchers (De Vany, 2004; Elberse & Eliashberg, 2003; Hennig-Thurau et al., 2007) to be influential in long term box office revenue. It is not sufficient to blast all information on a film to potential viewers in the first week, as that information decays and is drowned out by information on new films.
- *Environment*: We need to consider what else is going on in the environment to affect how a film will do. A film in a less competitive environment will outperform an identical film in an environment where there are other more successful running films or lots of new openings. Even a film making a certain amount of revenue by a certain age may survive in one environment but die in another more competitive one.

- *Information network structure*: Starting information in an area of the network where people don't attend films often, or don't have many neighbors, will result in slower diffusion of that information initially. Combined with 'time', this could serve to kill a film, where the identical film might survive much longer when starting in a different part of the network.
- *Heterogeneous agents*: every agent has a different level of awareness of each film, and attends films at a different frequency throughout the year. Each film's path to success will differ based on which agents hear about it, how much they hear and at what stage in the film's life they hear it, as well as the other environmental factors we have already mentioned.
- *Path dependency*: holding all of these other things constant, each film's performance will be greatly influenced by the starting conditions. In this case, where the information transmission begins in our scale-free network will influence how the information continues (or dies out) over time. Every time-step presents a new environment with a new set of parameters, and the end result will be different for each film.

In addition to these factors, the non-normal distribution of auto-correlated movie revenues forces us to apply modeling approaches that do not rely on the statistical assumptions necessary for unbiased parameter estimates. Both De Vany (2004) and Collins et al. (2002) posited that film revenues exhibit unbounded variance, which “undermines much of the existing work relating a film’s performance to its identifiable attributes within an OLS model.” (Collins et al., 2002). An unbounded variance violates the Law of Large Numbers (LLN) which is a vital assumption for estimating unbiased parameters through statistical modeling. The LLN states that the larger the sample size, the more we converge on the true mean, or the less likely we will draw a number that diverges significantly from what we have already drawn. In essence, the LLN determines the acceleration of convergence to the mean.

When a sample exhibits a Pareto distribution, this is an indication that the population has an unbounded variance, meaning that the variance grows with sample size. With this distribution, as $n \rightarrow \infty, \sigma^2 \rightarrow \infty$, thus it is considered unbounded or infinite. Think about drawing from a Gaussian vs. Pareto distribution: with equal sample sizes drawn from each, you are more likely to have drawn the largest values (large enough to significantly affect the mean and variance) from the Gaussian than you are from the Pareto distribution. In the case of movie revenue, the more sample films we draw from the population, the more likely that at least one of them will have revenue significantly greater than the ones we previously drew which will drive the mean higher, because only 20% of the movies become blockbusters and earn significantly more than the other 80%. The likelihood of drawing a blockbuster is much smaller, but has a more significant effect. The result is that convergence to the true mean is slowed, and in more extreme cases the sample may never converge to the true mean.

I chose NetLogo (Wilensky, 1999) to build ChatterMod due to its powerful yet intuitive nature, allowing me to focus more on the model itself and less on the programming and interface. NetLogo 1.0 was first introduced in 2002 by Uri Wilensky, Director for Northwestern University's Center for Connected Learning and Computer-based Modeling. Since that time many new capabilities have been added, including social networks and geospatial information systems. Scientific research incorporating NetLogo models has been published in numerous journals and conference proceedings across a wide range of disciplines; Table 3-1 depicts a small sample of these.

Table 3-1: Some of the many peer-reviewed journals and conference proceedings featuring research using models develop in NetLogo. Source: NetLogo References (2016). Retrieved from http://ccl.northwestern.edu/netlogo/references

International Journal for Parasitology
International Journal of Applied Earth Observation and Geoinformation
International Journal of Computers for Mathematical Learning
International Review for Spatial Planning and Sustainable Development
International Review of Economics Education
Journal of Business Research
Journal of Cognition and Culture
Journal of Disaster Research
Journal of Ecology
Journal of Economic Dynamics and Control
Journal of Environmental Radioactivity
Journal of Theoretical Biology
Proceedings of the ACM Symposium on Applied Computing
Proceedings of the 9th IEEE Int'l Conference on Industrial Informatics
Proceedings of the 11th IEEE Int'l Conference on Advanced Learning Technologies

NetLogo's intuitive language and approach not only assist the scientist developing agent based models, but also support the transparency and re-use of the model by making it easier for others to review, understand and apply it.

Figure 3-1 illustrates the flow of my model, with routines executed by the observer, persons or films:

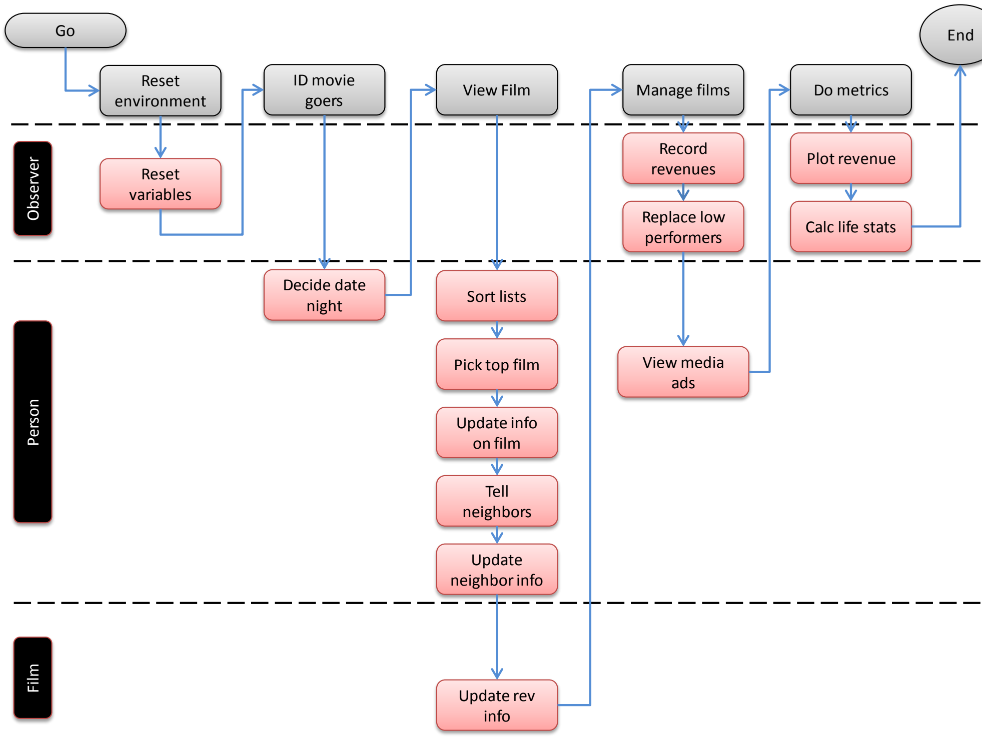

Figure 3-1: Conceptual diagram of the model flow

Following is a brief introduction to each routine, with more in-depth discussion in subsequent sections 3.2 and 3.3:

Reset the environment: There are variables that I reset at the start of every new time step. For persons, I reset their viewing statistics for the week so they are ready to consider viewing and discussing new films and trailers. I also reset variables related to time.

Identify movie-goers: Next I need to determine who will go see a film each week. Movie-goers are assigned frequencies when they are created which determine the likelihood that they will go see a film each week: *rarely*, *often* and *frequently*. With a probability corresponding to their frequency, movie-goers will decide if the current week is 'date night' and if so, will be designated to view films and trailers.

View films: Now the person is ready to see a film chosen from his *runningFilmList*, which is sorted based on volume of information known about the film. The person will choose the top film from the sorted list and go see that film, then tell his neighbors about it with some probability set by the user. *Revenue* is updated for the film based on the number of viewers that week.

Manage the films: Every week I need to perform some maintenance on the inventory of films. First, I send out advertisements of the pending films to the population of persons. Next, I calculate and record the ads that have been shown this week for each pending film, and the revenue earned during the week for the running films. Based on the revenue info, I close the underperforming films and replace them with the pending films.

Do metrics: Finally, I plot the revenue and calculate the life statistics of the films that have closed.

ChatterMod contains thirteen fixed parameters, nine of which are based on various literature sources as listed in Table 3-2.

Table 3-2: Model parameters

Parameter	Value	Basis
% who attend films frequently / rarely / often	0.162 vs. 0.147 vs. 0.691	Theatrical Market Statistics (MPAA, 2013)
Max films viewed per year	24 vs. 1 vs. 12	

Weekly decay rate of info	0.2	Advertising adstock literature (Joseph, 2006)
Likelihood of telling neighbors	0.7	Literature on word of mouth diffusion in a network (Allsop et al. 2007)
Advertising budget distribution	Exponential	Movie budget data (the-numbers.com, 2015)
Minimum running weeks	4	Industry exhibitor practice (De Vany, 2004)
Threshold for survival	Median	Calibrated based on data distribution
Word of mouth value	Varied	Abstract values representing impact of WOM and media ads; related to Bughin et al. (2010) concept of information 'equity'
Media ad value	Varied	
Population size	500	Abstraction: potential future GIS research
Count of films	15	

Each of these parameters is described in more depth throughout the following section. As you can see, most of the parameters are based in prior empirical analyses and literature except for word of mouth and media ad values, the size of the movie-goer population and the number of films. Word of mouth and media ad values are intended to be relative values to each other, and are the parameters that I varied for the experiments described in Chapter 4. They do not have an analogous real-world metric because they are subjective measures of a change in the relative significance of a piece of information, which is an abstract concept.

The numbers of movie-goers and films were chosen based on a variety of factors, and do not represent real-world films per capita, for a variety of reasons. First is that my model does not have a geospatial aspect to it: physically we are only representing one movie theater with fifteen films running, but abstractly we are representing the entire box office industry. A possible future extension to ChatterMod would be to recreate it using a geospatial platform or the GIS capability in NetLogo, with a more realistic box office infrastructure and movie-goer populations. For that, one would need to accurately represent the location of each theater, the number of screens and size of local population at each theater location, the number of films shown on each screen in each week, and the various release strategies employed by film distributors which can vary by regions and are often based on regional cultural factors.

As you can see, there is a much greater potential for error, and much more work involved in a more accurate geospatial representation. My work in this dissertation lays the foundation for future research in this area, so my goal here was to use the smallest movie-goer population possible in order to have the model run as quickly as possible, and to limit the potential of introducing additional errors to the model. In testing, I found that the ratio of population to films did have a fairly significant impact on the results, and I chose the ratio that provided the best results given that all of the other parameters were in line with real-world behavior.

From the thirteen fixed parameters and five routines we observe the following output values:

- Movie-goer network node degree plot and clustering coefficient
- Numbers and percentages of tickets purchased annually by movie-goers with frequency = rarely / often / frequently
- Percentages of movie-goers with frequency = rarely / often / frequently
- Average tickets purchased per movie-goer annually overall
- Average / minimum / maximum weeks of film runs

- Percentage of revenue held by top 20% of films
- Percentage of films achieving maximum revenue at open
- Gini coefficient and power law exponent for revenue distribution
- Percentages of films living past 7 and 10 weeks
- For each film: ID, run weeks, max revenue week, opening revenue, weekly revenue, total word of mouth, total ads, opening date

In the next sections I discuss the process through which this output is produced in more detail.

3.2 - Information flow

Remember from Chapter 2 that Eliashberg et al. (2000) implemented a Markov Chain model of consumer information and adoption behavior, called MOVIEMOD. In their model illustrated in Figure 2-1, consumers experienced states of undecided, considerer, rejecter, positive spreader, negative spreader and inactive. State transitions are triggered by advertising and word of mouth; probability of exposure to advertising depends on spending, and the probability of exposure to word of mouth depends on the number of spreaders, frequency of interaction and time-span of active spreading.

Information flow in ChatterMod can be represented by a similar diagram, with some notable differences illustrated in Figure 3-2 and explained below:

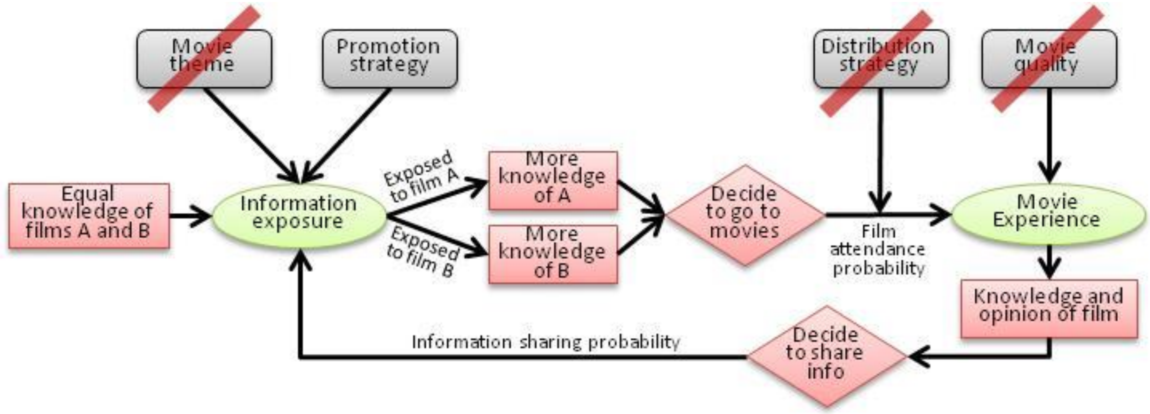

Figure 3-2: Information flow and agent states in ChatterMod

Figure 3-2 illustrates the flow of information in ChatterMod, simplified to a two-movie scenario to explain the basic premise. The agents in ChatterMod are in various information states as they gain information experience and movie experience, similarly to MOVIEMOD. However, the agents in ChatterMod are making a decision about which film to see relative to other films he knows about, unlike MOVIEMOD, where agents are making a decision about one film at a time without knowledge of other films. As with MOVIEMOD, ChatterMod has some parameters that the user can alter related to marketing and distribution, but unlike MOVIEMOD the user cannot alter move quality or theme since the movies in ChatterMod have no attributes.

In the simplified version of ChatterMod, consider a population of agents with equal knowledge of only two films A and B. Differing volumes of information are disseminated throughout the population for each film, such that after the dissemination each agent holds more information about one film than the other. Agents will then decide if it is 'movie night' – i.e. a good weekend to go see a film – and will then move on to experience the film. Afterwards, they'll disseminate information about the film to their neighbors with some probability.

The main difference between ChatterMod and the simplified version I just outlined is that ChatterMod has more than two films, plus several films waiting to open, and agents receive information on all of these films. Each agent maintains a ranked list of films in memory as shown in 'A' of Figure 3-3, and updates it at every time step with the new information he receives. One important thing to note is that, because each person has different information about each film, each person's list will have different values, and be sorted in different ways such that the top film for one person may not be top for another.

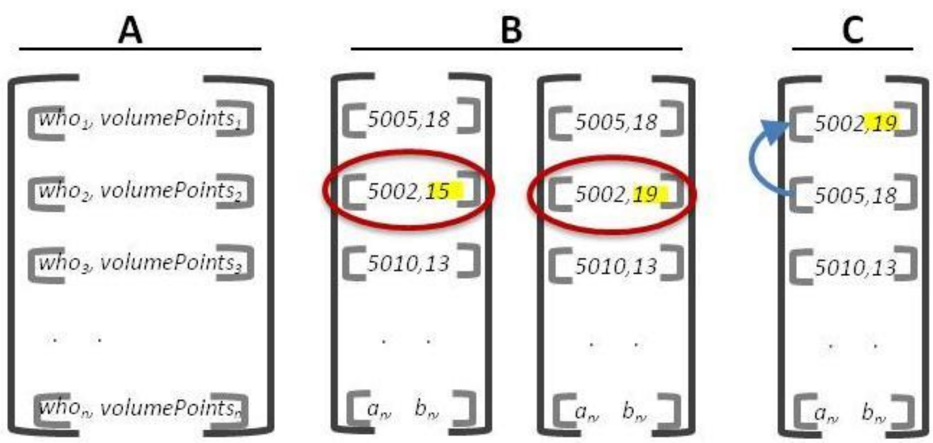

Figure 3-3: Updating and sorting of runningRankList

Figure 3-3 demonstrates an example of this in section 'B', where the person received information about film 5002 from neighbors and media ads. For each piece of information received, the person updates the *volumePoints* for film 5002, increasing that value from 15 to 19 in this example. After updating and prior to a person attending a film, their *runningRankList* is sorted by the amount of information the person has on each film, as in 'C' of Figure 3-3.

By sorting the list, I am putting the running films that are in the person's memory into the order in which the person will consider attending them, with film #1 being the first film the person considers. Since I sort by the volume of information the person receives on a film, the person will attend the film about which he heard the most. In the example here, the person will consider film 5002 first, as it has the highest volume score of 19. Remember that each person's list might be different because each person has different information about the films, so not everyone will have the same film at the top of their list. However, the more popular one film becomes, the more likely that film will be at the top of most people's list, since it will have generated the most buzz.

Once he attends the film he tells his neighbors about the film. Not everyone spreads this information; I set the likelihood at 70% based on research published by Allsop et al. (2007). This article presents some great detail with respect to information dissemination and adoption in a network, in fact making the distinction between those who provide information 'to a great extent' and those who provide it 'to some extent'. In my model, information is shared with equal likelihood by all, but since this is a stochastic process some will share more than others.

In the next sections I outline the model environment including the populations, their attributes and behaviors, the model routines and metrics output.

3.3 - Model routines and interface

The model interface is shown in Figure 3-4.

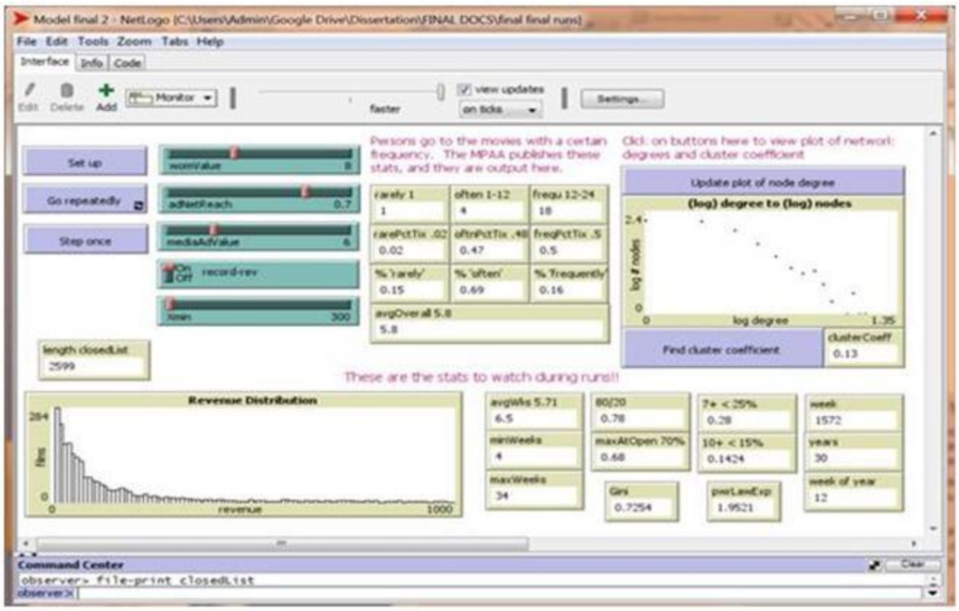

Figure 3-4: Model interface

ChatterMod has several parameters that are exposed to the user with the green sliders and chooser. There is also a number of monitors used for verification that the model output is in line with what we expect to see for real-life movie-goer behavior and film lifecycle.

Overview of populations and global variables

ChatterMod has two breeds of agents: *persons* representing movie-goers, and the *films* they attend. I set up these two breeds, and initiate the global variables, in the *setup* routine.

There are several global variables that are used to collect metrics on the model and how it is running. These can be classified as tracking time, revenue and other:

Time: Variables *week*, *year*, and *weekYr* track time steps of the model. The variables *life7Plus* and *life10Plus* are used to measure the proportion of films that last longer than 7 weeks and 10 weeks respectively. To calculate the maximum, minimum and average number of weeks for all film runs I use *maxWeeks*, *minWeeks*, and *avgWeeks*. Finally, the variables *rarelyCtYr, oftenCtYr, frequCtYr, overallAvgYr* and *frequPctYr* are used to calculate the number of times per year movie-goers attend films, for comparison to the MPAA statistics.

Revenue: These are primarily tables that hold the revenue, run and hype data on each film – *revTable, revAvgTable,* and *closedTable*. In addition I have the variables *top20Pct* which calculates the percentage of total revenue captured by the top 20 percent of films, and *maxOpenPct* used for calculating the percentage of films that earn their highest weekly revenue in their opening week. I also track the number of films that sell out – these are films for which all seats are filled.

Other: Finally, I have several variables that hold data on the structure of the network and other model characteristics. The variable *clusterCoeff* measures the cluster coefficient of the movie-goer network, and *degree* is used to calculate the distribution of degrees across all nodes. Finally, I use *closer* to represent the current film being closed due to underperformance.

Movie-goer attributes

There is a network of individuals who attend films, send and receive information on films via advertising and word of mouth. These movie-goers have no demographic attributes such as age or gender, however they are distinguished by the following attributes:

<u>Frequency of film attendance (*freq*):</u> At each time step, individuals will 'go to the movies' with a likelihood based on this frequency attribute, which is assigned to each person according to MPAA statistics on the frequency of movie attendance shown in Figure 3-5 (MPAA, 2013).

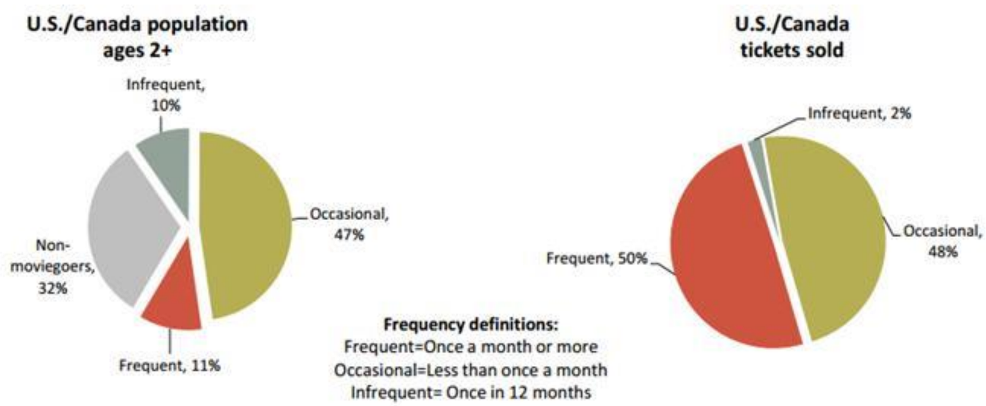

Figure 3-5: 2013 Movie-goer share of population and tickets sold. Source: MPAA. (2013). Motion Picture Association of America Theatrical Market Statistics

All of the persons in this network go to the movies at least once per year – I have made the assumption here that those who don't attend movies will have no movie-related information transmission activity, thus their existence in the network would have no effect on the box office. Therefore based on these statistics, 69.2% of the model's population will attend a film less than once per month, 16.2% will attend one or more times per month and 14.6% will attend about once per year.

In order to validate the movie-goer behavior of the model, I ensure adherence to the annual attendance frequencies above, as well as the average overall annual ticket sales rate of 5.8 tickets per year and the proportions of tickets sold as 50% to ‘frequent’ film goers, 48% to ‘occasional’ attendees and 2% to ‘infrequent’ film goers. This results in the following equation:

$$E(TicketsPerYear) = \frac{(avgOverall * pctTix)}{popPct}$$

We are given *avgOverall* from the MPAA as 5.8 tickets per year, *popPct* as 69.2% for ‘often’, 16.2% as ‘frequent’ and 14.6% as ‘rarely’ and *pctTix* as 50% to ‘frequent’, 48% to ‘often’ and 2% to ‘rarely’. Given those numbers, the expected tickets per year for ‘rarely’ is 1, for ‘often’ is 4 and for ‘frequently’ is 18: these all fall within the expected ranges of 1 for ‘rarely’, 1-12 for ‘often’ and 12-24 for ‘frequently’. I instruct the agents to choose a film if a randomly chosen number out of 52 falls below 1, 4 and 18 at each time step, and this procedure gives us the metrics that match the MPAA statistics.

<u>Ranked film lists (*runningRank and pendingRank*):</u> Each agent has two ranked lists of films of which they are aware, one that contains currently running films *runningRank* and another that lists films ‘coming to a theater near you’ *pendingRank*. Once a movie opens, it moves from *pendingRank* to *runningRank*. The ranking of these lists reflects the volume of information the agent receives on each film; it does not reflect any kind of preference data on the part of the agent. At each time step an agent chooses whether or not to see a film based on the frequency likelihood above; if he/she decides to

go to the movies then he/she will choose from the *runningRank* list and the film will subsequently be removed from the list.

Current movies, trailers and attendance (*myFilm*, *myTrailer*, *myFilmsRank*, and *movieNight?):* These are used to hold the current week's info for each film attendee. Every week, an attendee may go to see a movie with some probability, according to his assigned frequency, and this is stored as a Boolean 'true' or 'false' in *movieNight?*. The other variables hold the info on the film and trailer the person will see if *movieNight* is 'true'.

Movie-goer network

Schmidt et al. (2011) examined the effects of network structure on activity metrics by testing percolation and memory-less processes on three types of networks: an Erdos-Renyi random graph, a Watts and Strogatz small-world network, and a scale-free network. The percolation process in this study represents auto-correlation of node activation, and the memory-less process is of spiking activity where nodes fire at a Poisson rate. The authors found that the percolation process results reflected the degree distribution of the networks, while the memory-less process produced either exponential or heavy-tailed results independent of the network structure.

Each film's path to success or failure is just that - a path - and thus is definitely not a 'memory-less' process. De Vany (2004) found autocorrelation present in his study of movie revenues. Numerous authors found that word of mouth and/or media ads affect subsequent revenue (Burzynski & Baker, 1977; Elberse & Eliashberg, 2003; Hennig-Thurau et al., 2006; Liu, 2006; Mizerski, 1982; Sawhney & Eliashberg, 1996 to name a few). All of this indicates that the structure of the network should be as representative of the real world as possible, to ensure a more realistic information flow. Individuals relay information through their relationships with others, thus the relationships serve as channels for information.

These channels can be more personal as with friendship or family, or can be more virtual as with Facebook, LinkedIn or even YouTube. It is impossible to know all of the relations through which a single piece of movie information can travel, so I construct an artificial network as an abstraction of the real life movie-goer network without regard to the media type of the relationship. The key here is to ensure that a channel exists for information flow, not necessarily the specific type of channel.

Barabási and Albert (1999) identified the existence of scale-free networks arising naturally in many domains including the World Wide Web and social networks, and produced an algorithm for constructing scale free random networks using preferential attachment, which was named after the authors: the Barabási-Albert (BA) Model. The significance of the BA Model approach is that it generates a scale-free network – meaning the network exhibits a power-law degree distribution. The algorithm constructs the network using a mechanism called preferential attachment, where new nodes are more likely to connect to nodes with more links.

Interestingly, Barabási in conjunction with Ginestra Bianconi also identified the Bose-Einstein condensation process occurring as a phase in evolving scale free networks (Bianconi & Barabási, 2001); De Vany and Walls (1996) identified the Bose-Einstein process as occurring in the distribution of box office revenues. In their paper, Bianconi and Barabási (2001) identify three "thermodynamically distinct phases" of evolving complex networks as 'first-mover-advantage', 'fit-get-rich', and 'winner-take-all', the latter being the Bose-Einstein condensation process. It is this 'winner-take-all' phenomenon that De Vany used to describe the competition among films at the box office. An illustration of the relationship from their paper is in Figure 3-6:

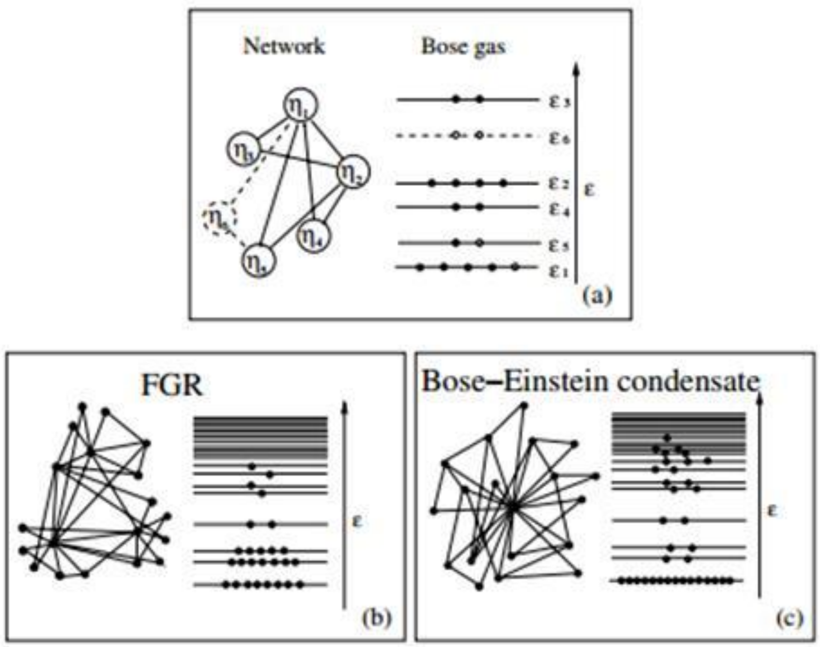

Figure 3-6: Mapping between network model and the Bose gas. Source: Bianconi, G., & Barabási, A.-L. (2001). Bose-Einstein Condensation in Complex Networks. Physical Review Letters, 86(24), 5632–5635

Bose-Einstein condensation describes the process of attraction of new entities (nodes, energy) to existing ones based on their fitness. This process in complex networks is measured by new nodes being attracted to higher degree existing nodes, in thermodynamics by particles being attracted to certain highly populated energy levels, and in box office revenues by a greater proportion of viewers being attracted to one film. The analogy to condensation is that these nodes / particles / viewers are *condensing* to one winner as the environment evolves, causing a shift from the 'fit-get-rich' phase to the 'winner-take-all' phase. In the case of Bose condensation, boson particles are increasingly attracted to one energy level as temperature falls.

The BA Model for generating scale free networks incorporates preferential attachment and network growth by basing the probability $\prod$ that a new node will be connected to an existing node i on that node's degree of connectivity k_i relative to the degree of all other nodes j: $\prod(k_i) = k_i / \sum_j k_j$. An indication that the network is scale free is when *P(k)* (the proportion of nodes having degree *k)* is $P(k) \sim k^{-\gamma}$ for $2 > \gamma > 3$.

Applying the BA Model for generation of the network in this study yielded $2 > \gamma > 3$ for $k > 6$ and $1.65 > \gamma > 2$ for $k =< 6$ in a network of 2000 nodes with k ranging from 1 to 33. The reason for the lower values of γ for $k =< 6$ may be explained by the addition of an adjustable clustering coefficient in ChatterMod. I incorporated this additional step in generation of this network that adds the ability to adjust the clustering coefficient during network generation, based on the research of Mislove et al. (2007). Mislove et al. (2007) observed that different social networks exhibit very different clustering coefficients not only from each other but from random graphs generated using preferential attachment. Their results are shown in Table 3-3:

Table 3-3: Clustering coefficients for various social networks, compared to generated networks of the same size. Source: Mislove et al. (2007). Measurement and Analysis of Online Social Networks. In IMC'07. San Diego, CA

Network	C	C / C Erdos-Renyi	C / C BA Model
Flickr	0.313	47.2	25.2
LiveJournal	0.330	119.0	17.8
Orkut	0.171	7.24	5.27
YouTube	0.136	36.9	69.4

The clustering coefficient measures the occurrence of 'triangles' in the network, or, the likelihood that two friends of A are also friends with each other, and this table demonstrates that the coefficient can vary widely among different social networks. In a social network of individuals with common interest (films) transmitting information on those films through various mechanisms, the coefficient is likely to be higher than perhaps the Web, but maybe not as high as Flickr. Film opinions generated by WOM are likely to be transmitted via several different mechanisms including direct friends and family, as well as online using Facebook or other sites, and we have no way of learning and modeling all transmissions so we must estimate. Because different clustering coefficients may be correct in different situations, I decided to add an adjustable clustering coefficient for future research.

The approach for incorporating an adjustable clustering coefficient was proposed by Herrera and Zufiria (2011) and is quite straightforward. With some probability, a new node will create a second link between itself and either a first- or second-degree neighbor of its first linked neighbor, as illustrated in Figure 3-7.

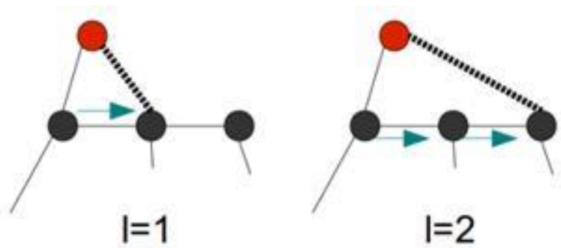

Figure 3-7: Creation of links for clustering coefficient. Source: Herrera, C., & Zufiria, P. J. (2011). Generating Scale-Free Networks with Adjustable Clustering Coefficient Via Random Walks. arXiv Preprint arXiv:1105.3347

Applying both the BA model for generating a scale-free network and the Herrera & Zufiria approach for ensuring a particular clustering coefficient, my network generation algorithm steps are as follows:

The *setup-populations* routine has three main steps: construct and link the first three nodes, construct the rest of the network, then format the persons. In order to set the parameters for the routine, the user chooses values for the maximum population size *maxPop*, probability that a new node will be assigned one neighbor or two (*singleProb*), and probability that the two neighbors assigned will be one or two degrees apart (*clusterPct*). Then, during the *setup* routine:

- Step 1: Create an initial set of three nodes connected to each other, each with degree $k = 2$
- Step 2: Randomly choose a value q such that $0 > q > \mu_k$ (where μ_k = mean degree of nodes)
- Step 3: Choose an existing node with $k => q$ and link new node to it
- Step 4: Randomly choose a value r such that $0 > r > 1$
- If $r <$ *singleProb*, then return to Step 1 to create a new node

- If $r > singleProb$, then proceed to Step 4 to add second link
- Step 5: Choose another random value *s* such that $0 > s > 1$
- If $s < clusterPct$, then create second link between the new node and the first-degree neighbor of the existing node
- If $s > clusterPct$, then create second link between the new node and the second-degree neighbor of the existing node
- Step 6: If count of nodes < *maxPop* then return to Step 1 and repeat, else end.
- Step 7: Finally, I format the persons, including setting the frequency with which they attend the films based on MPAA statistics.

Once *setup* is finished running, we can use a button on the interface to run the *calculate-pop-stats* routine to check the clustering coefficient and the log-log plot of node to degree. In the interest of reducing *setup* time, I do not calculate the degree and clustering statistics automatically.

Setting up the films

The films in ChatterMod are either pending opening, opened and running in theaters, or closed. These films do not have attributes relating to stars, genre, etc. and are in fact quite generic. The only attributes describing these films are related to the film's run, revenue and hype:

<u>Weekly status metrics *(runWeeks, revenue and status):*</u> Each week I collect statistics on the number of weeks the move has been running, the cumulative revenue which is counted as one per viewer and the film's current status which can be pending, running or closed.

Opening weekend *(openRev, openDate, openSeats and maxRevWeek)*: There are also statistics related to opening weekend that are of interest on their own, as well as being used in other calculations. These include the opening weekend revenue, the week number that the film opened, and the week that the film earned the highest revenue of its run.

Information metrics *(adsMedia):* Finally, each film will have a unique set of metrics that track all of the information activity associated with it. For instance, trailers for select openers (films waiting to be released) will be shown in theaters, and ads for those films will be broadcast to the general population via 'media'.

The main setup procedure for films – *setup-films* – has three subroutines. The first subroutine is *make-films* and is responsible for creating the initial population of film agents. The films being set up are both pending films (those that haven't opened yet) and running films (those that have opened). This subroutine calls others that will format the films, generate an initial revenue history for the running films and position them.

The formatting routine – *format-film* – is used for every film that is created during the simulation; these films will be pending films created to replace others that finally open.

Next I want to place the running films in their positions, and generate the revenue history. For this first inventory of films the revenue history is random, just so that the films have something in their history for the initial calculations. These films will die and their revenue totals will be disregarded, as will all films for the first 156 time steps (the equivalent of three years of movie-watching weeks), so these initial histories are of no significance. It is only after 156 time steps that the model begins to store the revenue statistics of each film.

Once the films are in place and their attributes are set, I need to populate the film lists of the persons. Each person carries two film lists: *runningRankList* and *pendingRankList*. These hold the running and pending films that exist, as well as the *volumePoints* relevant to each person on each film. All films go into each person's list, and again a random value is assigned to each film, this time for *volumePoints*. These values will differ for each person, so that each person has a different amount of information on each film. This heterogeneity with respect to information will hold true throughout the entire model run, reflecting the imperfect information that movie-goers in real life have on films.

Go **overview**

The go routine is repeated as weekly time-steps, stepping through the subroutines needed to choose the movie-goer, send him to the theater to view a film, have him spread information about what he saw, determine which films will be closed, open and close the file that holds individual film info, and calculate the metrics that are on the interface.

There are some variables that we need to reset each period in order to carry out those sub-routines. These are variables that are used to describe whether a person will go to a film (*movieNight?*), and what they will see (*myTrailer* and *myFilm*).

Remember that the movie-goers have been assigned classifications – *frequently*, *often* and *rarely* – designating how often they will attend films during a year. A percentage of the population receives each designation, and each attends a film each week with some probability corresponding to their classification. These frequencies and the percentage of population receiving them correspond to MPAA (2013) statistics: frequent movie-goers attend 12 to 24 times per year, those who go often attend one to 12 times and those who go only rarely attend about one film a year. I use a subroutine called *decide-date* to calculate and report the Boolean *movieNight?* back to the *ID-movie-goers* routine for each classification of movie-goer: a certain percentage of movie-goers will have *movieNight?* set to 'true' which tags them for film attendance this round.

Next, these tagged individuals must attend a film. To do this, we need to sort the person's film lists, choose which film they will see, update the film's revenue, and update the person's lists and viewing history. The first thing is to discount the older information a person has on currently running films. In marketing literature, the decline of an individual's memory of media ads is referred to as Adstock Decay; in ChatterMod this decline is in all of the information the person has on a film, be it from media ads or word of mouth. Many things affect the rate of decay, but in general the 'half-life' figures for TV and radio are 2-6 weeks and 1-5 weeks respectively. This means that, with a half-life of 3 weeks, a person's memory of a media ad will decline by half after 3 weeks. In ChatterMod, I have chosen an Adstock Decay half-life of about 3 weeks. What this means in practical terms is that the *volumePoints* value that a person has for a film will decay by 0.20 each week. Note that this occurs prior to a person getting new WOM and Media info in each time step.

Next, I have each person sort their film lists by the *volumePoints* representing the number of times the person has heard about each film. This means that the film about which the person has heard the most will be at the top of the list. Remember that this is a model of movie-goers with heterogeneous and imperfect information, such that not all movie-goers know about all films. Since every movie-goer has received information locally, not all movie-goers will have received information on every film, so it is entirely possible that a person will have zero *volumePoints* for one or more films in their list. This means that even though the film is in the person's list, they really have no idea the film actually exists.

Now it's time to choose the film. This subroutine starts by choosing the film that is at the top of the person's sorted list. The *who* of the film chosen by the subroutine becomes the person's *myFilm* value. Once the person views the film, I start updating the info of the film and person by adding one to the revenue of the film and having the person tell their neighbors about the film.

Once everyone who will see a film has done so, and revenues have been updated, I move on to managing the inventory of running films. This involves advertising the pending films, totalling all of the ads for each film, recording revenues, and updating the inventory of films. The first thing is to advertise the pending films; each film has a probability of being viewed by one or more movie-goers depending on the *adBudget* value randomly assigned when the film opened. Films with higher *adBudget* values will have more of a chance that someone will see them than those with lower values.

Then I record each film's revenue for the week, and update the inventory of films based on performance. The main routine for this is *change-films* and it calls two subroutines responsible for closing underperforming films and opening a new film from the pending inventory to replace the closed one. First, we need to close the films that are doing poorly. In reality, there are many aspects to this decision, mostly revolving around the desire to make room for other pending films for various reasons. Since many of these decisions involve intuition, incomplete information and best-guesses on the part of the exhibitor, I decided to keep the logic more simple and straightforward here, closing films that are achieving revenue below the median of all films for that week. It made the most sense to use the median rather than the mean, since the latter is more affected by extreme values and we are working with discrete data.

To close a film, I first set the film's *status* to *"closed"*, then record the max revenue week of the film. I then record the information on closed films that have revenue greater than zero, but only once the simulation has completed enough time steps to stabilize; I am confident of stabilization after 156 time steps (weeks), which equates to three years. Finally, I output the film's info to an external file for further analysis.

For every film we close, we need to open a new one. I choose the new running film from the inventory of pending films, and create a new pending film to replace it. For the opening film I record the opening week, and have all persons move the film from their *pendingRankList* to their *runningRankList*. Remember that while the pending film was waiting to open, people saw trailers for the film and have already accumulated information about it in their *pendingRankList*; this information will transfer to the *runningRankList* as well.

The final thing we need to do is to calculate the metrics and create the revenue plot on the interface. These steps are relatively straightforward. First, I plot the revenue as a histogram. Next, I update the time metrics. The model timestep is one week. As the model is running I track total weeks, total years, and week number within the year. As part of this step I reset a couple of values that are reset every week and year. The final subroutine calculates the statistics on films and persons. The films included in the statistics are all the films that have closed and were sent to the *closedList*.

——••••——

> *[Columbus]: Oh, America. I wish I could tell you that this was still America, but I've come to realize that you can't have a country without people. And there are no people here. No, my friends. This is now the United States of Zombieland.*

Quote from the movie *Zombieland*

CHAPTER 4 - MODEL VERIFICATION AND VALIDATION

Agent-based models incorporate many characteristics that make quantitative model validation much more difficult, and my model is no exception. There are primarily two issues with validation of agent-based models: they are usually modeling behavior of boundedly-rational humans in an open system, and thus the highly path-dependent result of this behavior is non-linear. (Cilliers, 1998). Simply put, this means that what happened in the past is not guaranteed to happen to the same magnitude in future, rendering point estimates unreliable as stylized facts. As an example, observe the degree to which the total tickets sold for the highest grossing film each year fluctuates from 1990 through 2015 in Figure 4-1.

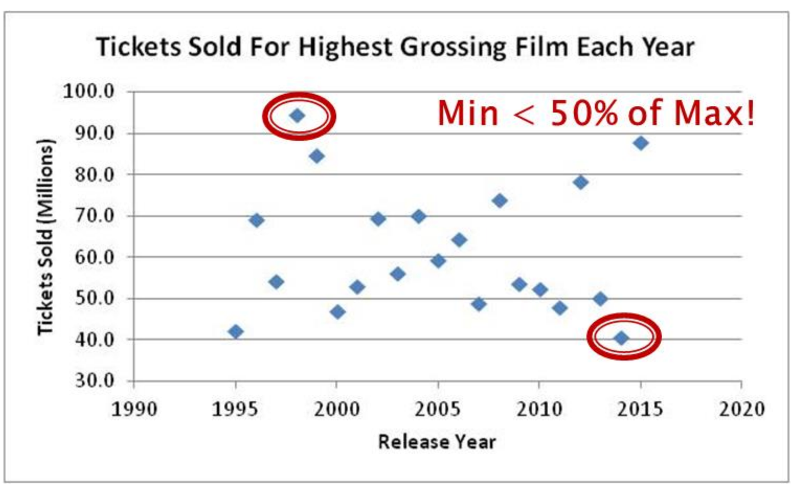

Figure 4-1: Number of tickets sold for highest grossing film each year. Source: "Top Movie of Each Year," 2016, www.the-numbers.com

Between 1995 and 2015, the lowest number (40M) of tickets sold for the top film is less than 50% of the highest number (95M) of tickets sold. Another example is the difference between the film statistics found by De Vany (2004) and those I found from a sample of more recent films listed in Box Office Mojo (2014), illustrated in Table 4-1.

Table 4-1: Comparison of film lifecycle statistics found by De Vany (2004) and more recently from Box Office Mojo (2014)

	De Vany (2004)	Box Office Mojo (2014)
Average film life (weeks)	5.71	6.6
% live past 7 years	25%	39%
% live past 10 years	15%	21%

In Chapter 2, I discussed the variation with respect to breadth (number of years of films included) and depth (top 50 or 100 vs. all, US only vs. international) in data sets used across the different studies of box office success. Imagine the variation inherent across those data sets with respect to estimates of tickets sold per year.

Strict adherence to such estimates would not be useful for testing a model that is not intended to make predictions but rather explore relationships. My model has been developed to explore the relationship between information dynamics and box office revenue, not predict what makes a particular movie successful or even what the total industry revenue will be. I want to know if the macro-level Pareto distributions of box office revenues we see in real-life can be generated by a model that represents the micro-level interactions of heterogeneous movie-goers who choose from a supply of homogenous films based solely on the level of imperfect information they have about the existence of each film.

Sargent (2011) stated that level of validation of a model should be decided with respect to the purpose of that model. My goal with ChatterMod is to grow the Pareto distribution of box office revenue from the ground up in order to describe the process by which this occurs, thus my approach to verification and validation focuses more heavily on the distribution of revenue and its relationship with information dynamics. Given this goal and the notion that stylized facts can change over time, my approach to validation against historical metrics is akin to that recommended by Popper (1959) who advocated for focusing more on the degree of model validity rather than a more absolute 'pass' or 'fail': I compare multiple metrics to prior empirical results from literature with the goal to align with as many as possible, but accepting that some misalignment does not imply total failure.

For ChatterMod I performed *verification* that the model was built and working correctly, and *validation* to determine the accuracy of the model to the extent appropriate given the discussion above. Sections 4.1 and 4.2 describe my verification and validation processes in more detail.

4.1 - Model verification

For verification I applied the trace testing approach: for each subroutine I studied the state of the system before and after execution. This section describes the tests I performed to ensure that the model has been constructed correctly, and is arranged according to each subroutine that was tested.

Setup population: The objective of this routine is to create a scale free network of movie-goers. Table 4-2 illustrates the test and actual values, showing that the model is producing the correct values.

Table 4-2: Testing - *Setup populations* test values

Parameter	Parameter Sources	Test Values	Model Output
Number of persons	User set	1000	1000
Movie-goer *frequency*	MPAA (2013)	'Often' = 69.2%, 'Frequent' = 16.2%, 'Rarely' = 14.6%	'Often' = 69%, 'Frequent' = 16%, 'Rarely' = 15%
Clustering coefficient	Mislove et al. 2007, for YouTube	c = 0.136	c = 0.14
Node degree distribution	Scale-free networks have: $P(k) \sim k^{-\gamma}$	$2 > \gamma > 3$	$\gamma = 2.26011$

Setup films: This routine creates two sets of films; one set is currently running and the other is a set of pending films. With this routine I also want to assign random starting histories to each of the films. This test is simply to ensure that the population of films represents what I expect to have; it does not represent a real world statistic.

In the test case, I asked for 20 running films and 20 pending films. I was able to confirm both of these settings via monitors on the model interface.

Populate person lists: Once the films and persons are created, we need to populate the persons' running and pending movie lists with film info. This routine populates these two lists for each person; each list contains the ID number of each film in both the running and pending inventories, and initial random values for information points, rating, and number of weeks. Upon completion of this task, I examined movie-goer lists of running and pending films to ensure they were correctly populated.

Reset environment: This routine resets a few of the person variables prior to each time step. In particular, the routine resets *myTrailer, myFilm, and myFilmsRank* to zero, and *movieNight?* to false. To test this, I took a count of '*movieNight?* = true' and a sum of all the other variables over all persons; both the sum and count were zero as they should be.

ID Movie-Goers: Next I identify the population of persons who will be viewing films this week. Remember that we have movie-goers with three types of frequency – rarely, often, and frequently – and annually they attend films once, 1 – 12 times, and 12 – 24 times respectively. These statistics are from the MPAA statistics on the frequency of movie attendance shown in Figure 3-5.

I was able to verify via the model interface that the population of 'rarely', 'often' and 'frequent' movie-goers has gone once, four times and 18 times per year respectively. I was also able to verify the percentage of tickets that each type of movie-goer has purchased during the year: *rarePctTix* value is 0.02 vs. metric of 0.02, *oftnPctTix* is 0.47 vs. metric of 0.48 and *freqPctTix* is 0.5 vs. metric of 0.5. Finally, the average tickets per year purchased by all movie-goers of 5.8 matches the metric of 5.8.

Decay information, sort lists and choose film: This routine manipulates the running films lists of each movie-goer chosen for 'date night' and chooses the film they will see. First, we need to ensure that the most recent information the movie-goers receive is weighted more heavily than older information. It is a simple calculation that reduces the information points on each film in a person's list by a user-set level; for this test, I set the level at 25%. Next, I sort the running rank list in order of information points, such that the film with the highest number of information points is at the top of the list. Finally I choose the film with the most information points and available seats.

To verify this, I examined the running rank list of one movie-goer before decay and after. Each film in the list has two values: the first film in the person's list has values 1338 and 1.83 representing film ID and information points prior to this routine. Afterwards, I verified the information points had fallen by 25% for film 1338, from 1.83 to 1.46. Note that movie-goers will also receive new information on each of these films during the time step, so the overall change to information points once the time step ends will not always be the same percentage decrease. If the film takes off, the information points may increase.

Also, the person’s running rank list had been sorted by information points: prior to the sort the first film in the list was film 1338 but after the sort the film moved to position #4.

View film (update film revenue): Once a person views a film, that film’s revenue increases by 10. First I verified that at the start of each week the revenue is zeroed out for each film, then I verified that films accrued new revenue for the week.

Update movie-goer and neighbor film info (tell neighbors): Once the movie-goer has viewed a film, she tells her neighbors about the film. For this step, I verified that a person's information points increased by the correct amount if they were told of a film by their neighbor.

Advertise Trailers: As with new films in real life, pending films are marketed using advertising. These trailers are shown broadcast via mass media, and the pending list of the person who viewed the trailer should be updated. My verification here is to make sure that the pending films in the person lists are indeed being updated with the additional information points. For this test, each viewing of a trailer results in an increase of six information points, so I verified that a person's list was updated with this change upon viewing a trailer.

4.2 - Model validation

In this section I move on to validating that the model produces the correct film lifecycle and revenue values for multiple metrics, under a range of values for word of mouth (WOM) and media ads. The approach I chose for validation is in line with the multi-stage approach recommended by Naylor and Finger (1967), which is to:

- Develop the model's assumptions and parameters based on all available information
- Where possible, test those assumptions and parameters for validity
- Test the model's ability to predict (replicate) the (historical) behavior under study

In Chapter 3, I described the construction of the model including the available information I used to construct it. Most of the data I used was emprically derived from prior studies, however I did also perform independent testing on several of the film statistics and compare them to De Vany (2004) as shown in Table 4-1. Finally, I tested the model's ability to replicate the distribution of film revenues, along with several other film lifecycle metrics, as in step 3 of the approach by Naylor and Finger (1967). The metrics and values for the model are outlined in Table 4-3.

Table 4-3: Sensitivity test values

Metric	Real-world values	
	De Vany[1]	**2015**[2]
Avg life (weeks) of film	5.7	6.6
% films living 7+ wks	< 25%	< 39%
% films living 10+ wks	<15%	< 21%
% revenue by top 20%	80%	80%
% earn max rev in wk 1	70%	71%

The average number of weeks, likelihoods of living past 7 and 10 weeks, and percent of films earning their max revenue in the first week are all statistics that come from De Vany's empirical analyses of box office revenue. The Gini coefficient is derived from a distribution where 80% of the revenue goes to ~20% of the films. And the Power Law exponent for infinite variance comes from statistical literature, while the fact that movie revenues display infinite variance has been found by De Vany (2004) and others (Sinha & Raghavendra, 2004).

Given the variation in point estimates over time that I discussed at the beginning of this chapter, the purpose of the model (to explore relationships rather than predict future point estimates), Sargent's (2011) assertion that level of validation should be determined based on purpose, and Popper's (1959) recommendation that the degree of validity rather than 'pass' or 'fail' should be the goal, my goal with this validation is to find the set of parameters that delivers the most reliable results, and results that could *realistically* be observed in such a system. For instance, De Vany (2004) observed that the likelihood of a film exceeding 10 weeks is less than 15%, and less than 25% for seven weeks. My subsequent analysis of films from 2015 found values of 21% and 39% respectively. So, it would be *realistic* to see simulation values that fall between those values, but it would not be *realistic* to see 0% of films reaching seven or ten weeks, or the majority of films reaching those life spans. While my validation process includes hypothesis testing for several variables, given the recent discussion of goals and point estimates, the failure to reject the test for every metric is not critical to acceptance of the model.

For this experiment I varied the values for word of mouth from one to 20 and varied media ad values also from one to 20, both in increments of one. This resulted in a total of 400 different parameter specifications being tested. I ran each parameter specification 50 times, for a total of 20,000 runs. Each time-step of the model is equivalent to one week, and each run lasted for 312 weeks or six years. Calculations did not begin until after the first three years, allowing for the model to stabilize so that the most accurate results could be produced. The output consisted of macro-level values for the metrics in Table 4-3, plus revenue and lifecycle values for each individual film.

The experiment starts at WOM value 1 and runs through media values one through 20 before changing to WOM value 2 and running through all media values again. The x-axis is numbered according to the configuration numbers, so configuration number 1 would be WOM value 1 and media value 1, configuration number 2 would be WOM value 1 and media value 2, and so on through configuration number 400. Table 4-4 depicts a summary of this numbering scheme.

Table 4-4: Explanation of configuration numbering

Configuration	1 - 20	21 - 40	41 - 60	...	381 - 400
WOM value	1	2	3	...	20
Media value	1 - 20	1 - 20	1 - 20	...	1 - 20

The initial results of the parameter sweep are depicted in Figure 4-2.

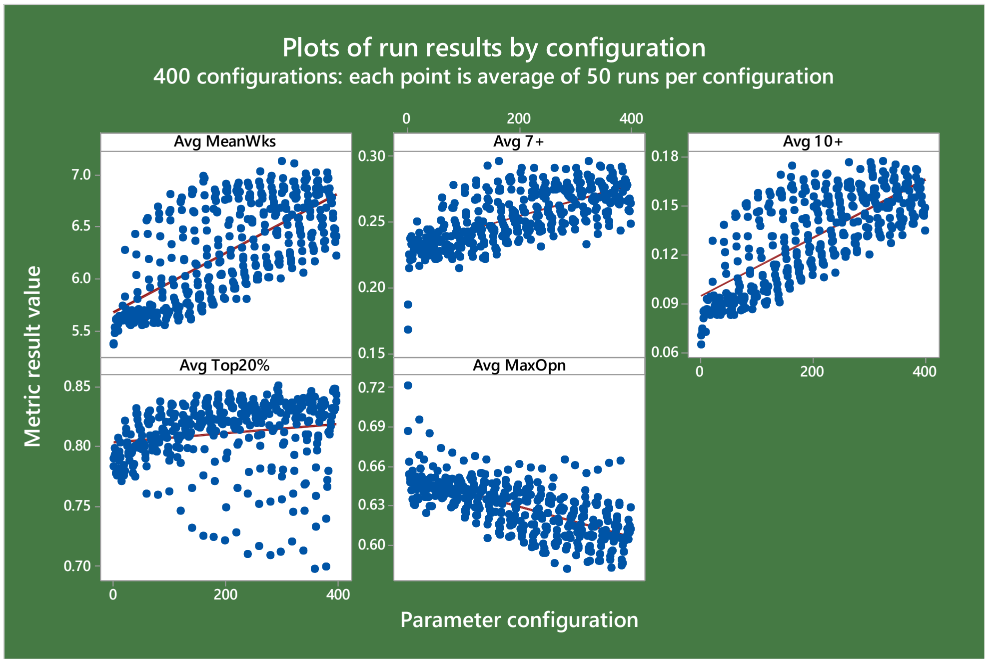

Figure 4-2: Initial sensitivity test results

Figure 4-2 shows scatter plots for each output, with one dot representing the average of the 50 runs for each configuration. Overall, with such a wide range of results for each set of configurations, it is difficult to tell what the 'best' parameter values are, particularly since the values improve as the configuration number increases for some outputs but decline for others.

One thing that does immediately stand out in these graphs, is that some runs exhibit more variability than others. Since one of the goals of validation was to identify the set of WOM and media values that provided more stable (i.e. repeatable) results, I calculated the Coefficient of Variation (CV) for the run results, defined as the ratio of the standard deviation to the mean. The benefit of using the CV in this situation is that it is dimensionless, so we can compare the variability of two or more samples with different means and even different units. It is commonly used in clinical research labs (Schechtman, 2013) as a reliability test for diagnostic tests and biochemical assays: in my case I am testing the reliability of my agent-based model for each parameter specification.

Figure 4-3 shows a graphical example of the variation in CV for the average lifespan of films (in weeks) for three different parameter configurations. The graph shows that run #93 has a CV of 0.031 and a mean of 5.453, run #210 has a CV of 0.042 with a mean of 6.273, and run #303 has a CV of 0.106 with a mean of 7.121. The most stable parameter configuration is the one with the lowest CV, so with this goal in mind run #93 is more desirable than the other two. Also of note is that the run with the smallest CV also has a mean closest to the target value of 5.71.

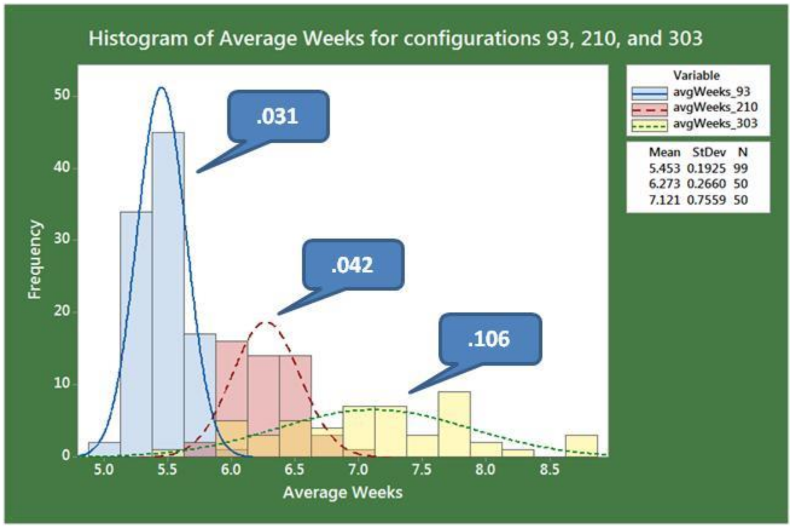

Figure 4-3: Example of variability in Coefficient of Variation for one output

Remember from Table 4-3 that there are multiple real world values to which I am trying to align model outputs. Figure 4-2 illustrates the difficulty in performing this task graphically: increasing WOM and/or media values decreases the variation of some outputs and increases it for others. The approach I chose was to sum the CVs over all outputs, and seek to minimize the *overall* variation of each parameter configuration.

Because the intended use of ChatterMod is as a way to explore the effects of word of mouth, media ads and network structure on box office statistics, it is not enough to identify one 'best' parameter configuration. The most useful information would be a range of parameter values across which experiments could be performed with *reliable* - i.e. low coefficient of variation - results. Logistically, it would be best to identify a contiguous range of parameters, to make experiment design easier.

With these goals in mind, I arranged the *overall* CVs (sum over all outputs) for each parameter configuration into a matrix, with the media parameter values on the x-axis and the WOM parameter values on the y-axis, as in Figure 4-4.

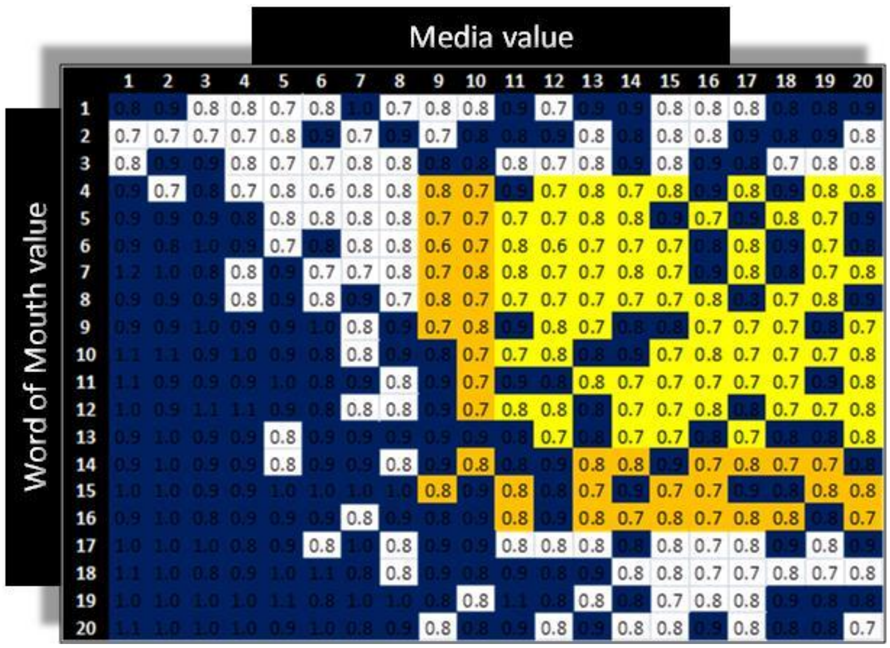

Figure 4-4: Coefficient of variation matrix for all experiment runs

To make it easier to visually identify the ranges of parameter values that minimize the sum of CV over all outputs, I colored cells dark blue conditional on ranking in the top 50%. I then highlighted in yellow a large contiguous range of cells, most of which were not in the top 50%, and highlighted an additional perimeter range in orange that could also be considered. The yellow-highlighted range - consisting of word of mouth values 4 through 16, and media ad values 11 through 20 - represents a range of parameters useful for experimentation because they deliver the most reliable results. However, we still want to determine how *accurate* the results from those ranges are, with respect to the real-world metrics.

Given the earlier discussion regarding the goal of ChatterMod and the variability in historical point estimates, my criteria for validation of ChatterMod was:

- First, identify one 'best' parameter configuration from the chosen range, where the majority of outputs aligned with the real-world metrics at the 95% confidence level
- Then, ensure that the average of most outputs over all runs in the chosen range aligned with the real-world metrics within the inter-quartile ranges.

Figure 4-5 illustrates the results for run #260, with WOM = 13 and media = 20; the red dotted line indicates the real-world test value. This run produced six outputs where the test value from De Vany's (2004) research fell within the 95% confidence interval of the results. The confidence interval for *life10Plus* - the percentage of films that live past 10 weeks - is between 0.11 and 0.13, which is still realistically close to the 0.15 test value even thought it failed the hypothesis test. Likewise for *maxOpenPct* - the percentage of films that achieve their maximum revenue in opening week - where the interval of 0.62 to 0.65 is still realistically close to the 0.70 test value. In fact, 35 of the parameter configurations produced at least four results aligning with test values at the 95% confidence level. One reason for the lower figures in my run compared to De Vany (2004) is that his analysis focused on films in the top 50, which means his figures would be biased upward from a study focused on the entire population of films. Since my model produces a population of films that are both 'winners' and 'losers' and my analysis includes all of those films, it is not surprising that my figures for lifespan may be lower.

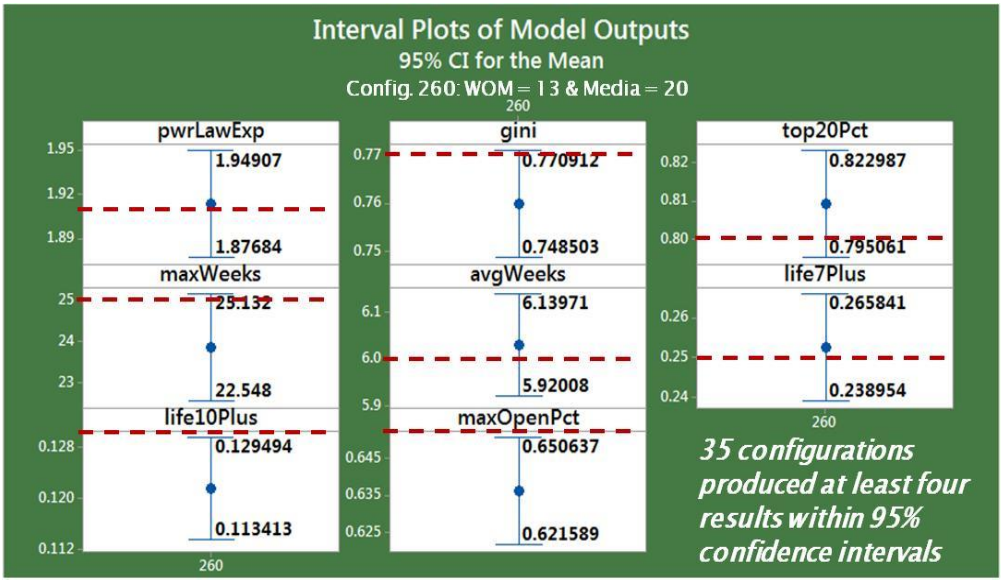

Figure 4-5: Hypothesis testing of run #260

For the second test, I wanted to see if the average over all parameter configurations produced inter-quartile ranges containing the real-world test values. Figure 4-6 illustrates the results of this test.

Of the eight outputs shown in this set of boxplots, five of them produced inter-quartile ranges containing the real-world test value. Again, *life10Plus* and *maxOpenPct* produced ranges slightly below the test value but still not unrealistic. The third parameter - *MaxWks* - was also low but realistic; again, given that my analysis focuses on both winners and losers (i.e. those who would and would not make the top 50, respectively) it is not surprising that my lifespan figures are lower.

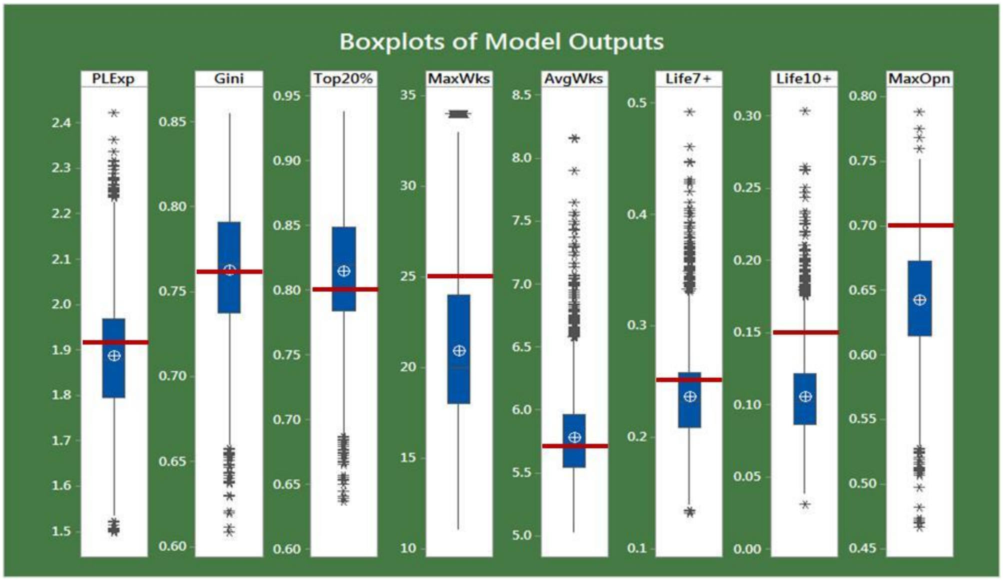

Figure 4-6: Boxplots of all model results within chosen parameter ranges

Thus far, I examined how the outputs of the model correspond to the real-world point estimates obtained through empirical research. Next I examined the input-output relationships of the model, specifically, the direction of change in output as WOM and media are varied. The results overall were consistent with those found in prior empirical studies in that word of mouth increased the lifespans of films while media ads were more influential in the opening week. Both the studies from Elberse and Eliashberg (2003) and Hennig-Thurau et al. (2007) found that ads were key for short term box office success while WOM was key for long term success.

Using the subset of output from the previous experiments pertaining specifically to the chosen range of parameter values (4 <= *womValue* <= 16 and 11 <= *mediaValue* <= 20), we can see how increasing the value of WOM actually increases the average survival of a film, while increasing media ad value decreases the survival rate. The relationships are shown in Figure 4-7.

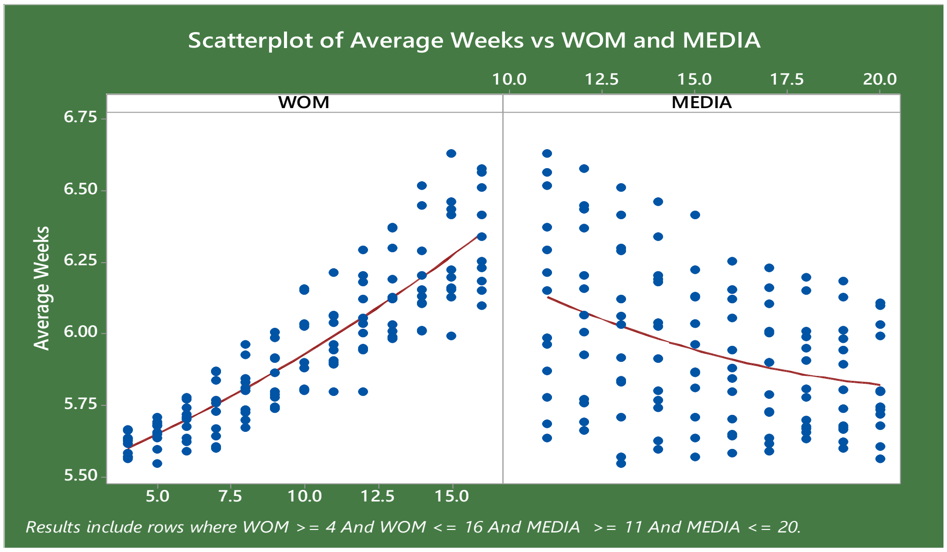

Figure 4-7: WOM and Media vs. avgWks

Remember from Chapter 3 that increasing the value of WOM or media means that each piece of information carries more weight - resulting in more value added to the total information associated with that film - and that people attend the film for which they have the most information.

These results are supported by the literature: Liu (2006) used social media data and found that WOM volume had a statistically significant effect on success but valance did not. Basuroy et al. (2003) found that valance of reviews did have an effect, but only positive buzz had a lasting effect. Eliashberg and Shugan (1997) saw reviews as predictors rather than influencers, and only on long-term box office revenues; they found that media ads had more of a relationship with short term box office. Sawhney and Eliashberg (1996) found that more heavily advertised films had exponentially-declining revenues while those that relied more on WOM declined more gradually. Elberse and Eliashberg (2003) found that advertising was key for opening but WOM was key for subsequent weeks. Gemser et al. (2007) also found that the volume of information was significant but valance was not.

De Vany (2004) defined success not only in monetary terms (films that achieved $50M or more) but also as films that gain 'legs' living longer than their competitors. With that in mind, the fact that *womValue* in ChatterMod positively affects the average life of a film makes sense, and is empirically supported. What is more interesting, however, is that *mediaValue* negatively affects average life. While Sawhney and Eliashberg (1996) did find that more heavily advertised films did have more quickly declining revenue, the implication was that this was related to the quality (or lack thereof) of the film. Since we have nothing that relates to quality in ChatterMod, we must consider other more 'mechanical' reasons for this effect. The most likely is that heavily advertised films in ChatterMod create more competition for movie-goer attention, such that when a film opens and moves from the person's *pendingRankList* to their *runningRankList* it is entering the list with an already high information score relative to the other existing films in the list. Such a phenomenon will affect the opening weeks of a film, but since advertising stops when a film opens, and memories decay in the model, the effect of the media ads will not affect the film's success in the long term.

These model results do support the empirical evidence that media ads have greater effect on big openings, and produce some interesting results with respect to WOM. One of the outputs of my model is *maxOpenPct*, which measures what percentage of films attain their maximum revenue week in their opening week. The output for this variable is plotted in Figure 4-8 against both *womValue* and *mediaAdValue*.

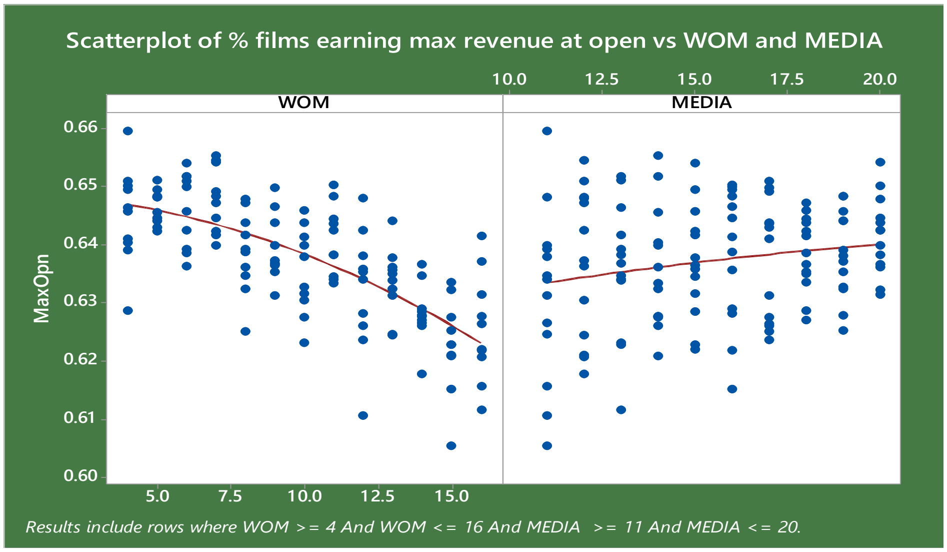

Figure 4-8: Scatter plot of maxOpenPct against womValue and mediaAdValue

Interestingly, WOM decreases the percentage of films earning their max revenue at opening. This indicates that in my model WOM gives a bounce to films that may start of slowly but gain popularity as word gets out about them. Also interesting is that increasing media ad value does increase opening revenue in my model, but not significantly. The effects of WOM and media on the pattern of revenue at opening vs. subsequent weeks in my model are supported by the study of Sawhney and Eliashberg (1996) who found empirically that film adoption followed three main patterns, depicted in Figure 4-9.

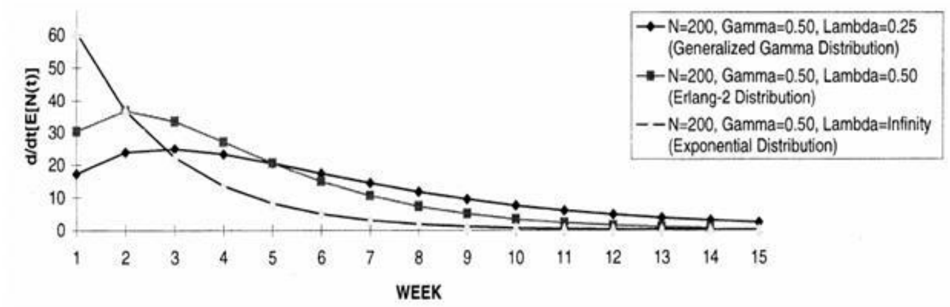

Figure 4-9: Sawhney and Eliashberg 'time to adopt' patterns. Source: Sawhney, M. S., & Eliashberg, J. (1996). A Parsimonious Model for Forecasting Gross Box-Office Revenues of Motion Pictures. Marketing Science, 15, 113–131

They noted that the films that are more heavily advertised tend to follow the exponential distribution, while the ones that survive more on WOM tend to follow the generalized gamma or Erlang-2 distributions. These latter distributions both increase at first, then decrease more slowly, which represent films that do not attain their maximum revenue in the first week.

The discussion so far has focused on how WOM and media affect the average life and opening of all films in ChatterMod. Next I turn to the subject of this thesis: the Pareto distribution of revenue across films and how it is affected by WOM and media. Figure 4-10 illustrates the effect that WOM and media have on the share of revenue earned by the top 20% of films.

We can see that increasing WOM value increases the revenue inequality among films, meaning that the top 20% of films earn a higher percentage of the total revenue, while increasing media value has the opposite effect.

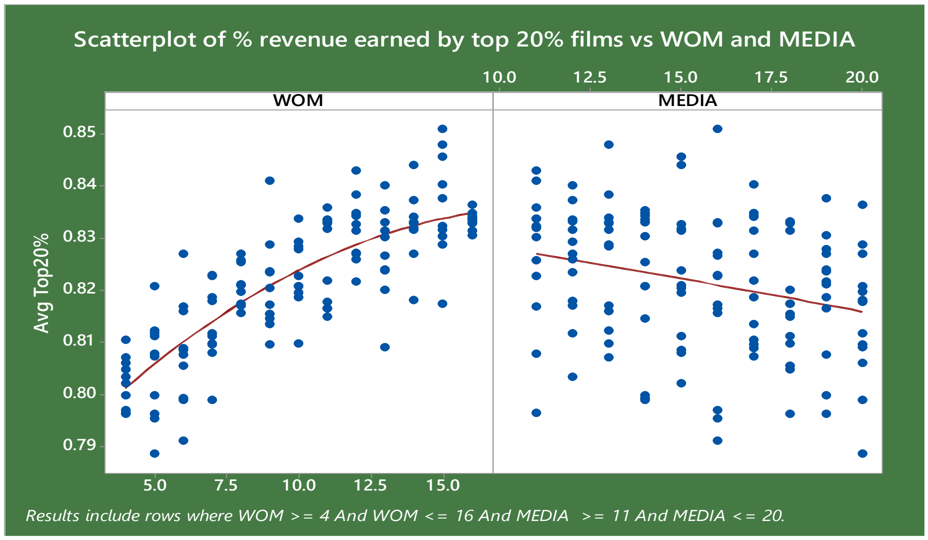

Figure 4-10: Distribution of revenue across films by WOM and Media

While increasing WOM value increases the average lifespan of films overall, it does not ensure that *all* films are successful; rather, it increases the opportunity for some to break away from the pack and become blockbusters. This finding is significant in that it supports the finding of De Vany (2004) that the Pareto distribution of revenue arises from the information dynamics among movie-goers. It is also significant because it shows that I have succeeded in what I set out to do here, which was grow this Pareto distribution from the micro-level interactions - sharing of WOM - among movie-goers.

In Chapter 5, I summarize my work and discuss future research possibilities.

——••••——

[Fin Shepard]: I know you're scared. I'm scared too. They're sharks. They're scary. No one wants to get eaten. But I've been eaten. And I'm here to tell you it takes a lot more than to bring a good man down. A lot more that to bring a New Yorker down.

Quote from the movie *Sharknado 2: The Second One*

CHAPTER 5 - DISCUSSION AND FUTURE WORK

De Vany (2004), as well as Sinha and Raghavendra (2004), observed this Pareto distribution of box office revenues, where 20% of the films earn 80% of the revenue. This finding renders attempts at predicting revenue for any single film - as well as the use of film attributes for attempting such a prediction - futile. When a sample exhibits a Pareto distribution, this is an indication that the population has an unbounded variance, meaning that the variance grows with sample size. With this distribution, as $n \to \infty, \sigma^2 \to \infty$, thus it is considered unbounded or infinite.

Think about drawing from a Gaussian vs. Pareto distribution: with equal sample sizes drawn from each, you are more likely to have drawn the largest values (large enough to significantly affect the mean and variance) from the Gaussian than you are from the Pareto distribution. In the case of movie revenue, the more sample films we draw from the population, the more likely that at least one of them will have revenue significantly greater than the ones we previously drew which will drive the mean higher, because only 20% of the movies become blockbusters and earn significantly more than the other 80%. The likelihood of drawing a blockbuster is much smaller, but has a more significant effect. The result is that convergence to the true mean is slowed, and in more extreme cases the sample may never converge to the true mean.

De Vany theorized that it was the complex information dynamics among movie-goers that produced such a statistical distribution. Several other researchers have found empirical evidence supporting the volume of buzz (Basuroy et al. 2003, Eliashberg and Shugan 1997, Sawhney and Eliashberg 1996, Elberse and Eliashberg 2003, Liu 2006, Gemser et al. 2007), with some finding that the valance of buzz did not matter. These researchers also found that media ads had a greater effect on opening and short term box office, but that WOM had a greater effect on long term box office, even enabling a film to pick up steam and earn its max weekly revenue later in its life rather than at open.

So, what is the driver behind the Pareto distribution of box office revenues we observe in real life and in ChatterMod? Previous research indicates that it is the percolation process of movie information throughout a scale-free network. Schmidt et al. (2011) examined the effects of network structure on activity metrics by testing percolation and memoryless processes on three types of networks: an Erdos-Renyi random graph, a Watts and Strogatz small-world network, and a scale-free network. The percolation process in this study represents auto-correlation of node activation, and the memoryless process is of spiking activity where nodes fire at a Poisson rate. The authors found that the percolation process results reflected the degree distribution of the networks, while the memoryless process produced either exponential or heavy-tailed results independent of the network structure. Several researchers found auto-correlation in movie revenues through WOM from previous viewers (Burzynski & Baker, 1977; De Vany, 2004; Elberse & Eliashberg, 2003; Hennig-Thurau et al., 2006; Liu, 2006; Mizerski, 1982), and based on the work of Schmidt et al. (2011) it is this auto-correlation percolating through a scale-free network that produces the Pareto distribution of box office revenues we observe.

I set out to organically grow this Pareto distribution of movie revenues from the micro-level information exchange among movie-goers, about a homogenous set of films and the results exhibited in Chapter 4 prove that I have achieved this goal. This is significant in several respects:

- It is the first agent-based model of the movie industry
- Unlike so many of the past models of the movie industry in the research described in Chapter 2, ChatterMod focuses purely on the information and behavior of movie-goers, treating the films themselves as homogenous
- Numerous outputs from the model correspond well to real world point estimates and distributions
- The agents in ChatterMod comprise a synthetically grown scale-free social network, and the information dissemination that drives movie revenue distribution occurs along the links of the network

While ChatterMod might deliver some bad news to those hoping to predict the success of a particular film, the aspects outlined above mean that ChatterMod is a significant leap forward in understanding how information dynamics affect the industry overall. Those with a more strategic interest in the movie industry will find that ChatterMod can aid in identifying more accurate real-world indicators of movie buzz and potential success. Advertising executives may employ ChatterMod to better inform strategies for generating word of mouth, shifting from opening-week media blitzes to more subtle, targeted and longer-term campaigns. Specifically, the use of a scale-free network to represent movie-goers can help executives identify the characteristics that make some movie-goers better targets for WOM-generating activities than others. Finally, ChatterMod's results can better inform the decision of the movie exhibitor, who is trying to determine (based on a variety of factors including level of buzz) whether or not to renew a film for another few weeks.

In addition to the potential employment of my model by movie industry and advertising executives, the model also provides fertile ground for future research in information dynamics, both general and specific to the movie industry.

With ChatterMod, I have provided a platform for future research in a variety of areas, including:

- Movie-goer Zero: More in-depth study into how the network location of the first movie-goer to hear about a film, as well as the path that the WOM takes through the network, determines the success of that film. One possibility might be to assign an 'early adopter' attribute to certain nodes such that they are more likely to be the first to view a new film, then study how success for the film changes as the assigned nodes change.
- Neighbor trust: In ChatterMod, I did not model the trust that a neighbor has in the information he is receiving, and thus his likelihood of acting on that information. I omitted this mainly because I are not conveying opinions about movies, just knowledge of the existence of the film. The addition of a trust attribute could be

added to adjust the value of the WOM or media ads, multiplying the info value by a trust value ranging from zero for no trust to one for complete trust.

- Information cascade: Further analysis of the model in its current state could be done to identify the points at which a film takes off and becomes destined for the long tail of the distribution. Such a study would examine how the information traveled across the network from both space and time perspectives, to see if there were any commonalities among such phase-shift events.

- WOM proxies: one could use ChatterMod to help identify real world proxies for WOM. Remember that Elberse and Eliashberg (2003) relied on proxies such as prior week movie attendance for WOM, which did not account for the social network connections of each of the prior attendees and the varying levels of information of potential future attendees. Liu (2006) was able to identify the best available sources of WOM data in social media, but this required manual classification of comments which meant only a small sample could be used. An approach would be to compare the WOM values and results of ChatterMod with Liu's results, in order to develop a model that can make use of actual available social media data in a scale-free network in a more automated way that doesn't require manual classification.

- Competitive landscape: back in Chapter 1, I discussed the idea that identical films can likely do very differently in different competitive environments. An interesting study made possible by ChatterMod would be to test that hypothesis by studying how one movie performs at the box office in environments where more competition exists vs. less competition. For instance, choose one movie that earns 500 word of mouth events (is talked about 500 times during its life) from each of 1,000 runs and characterize the competitive landscape by the WOM and rank of the other movies that performed better, to determine if any consistent patterns exist.

The research possibilities above can be done with little-to-no changes to the current model. One potential extension of my work requiring a larger modification would be to generate other types of graphs - e.g. the Erdos-Renyi random graph and Watts and Strogatz small-world network that Schmidt et al. (2011) used - for ChatterMod and examine how the results are affected by different network structures. This study would focus on whether or not a different result is obtained by incorporating a different network structure. If so, the next natural step would be to identify the real-world network structure of movie-goers and compare that to what is in the model.

To study the practice of phased movie releases across the US, one might consider incorporating a geospatial aspect to ChatterMod. In reality, movie openings are often phased across the country, for instance starting out in Los Angeles and New York and then opening in the rest of the locations based on performance. Such a release strategy would undoubtedly have interesting effects on the transmission of word of mouth, especially when it occurs via social media. A further extension along these lines would be to add international box office activity as well.

Finally, the model could be adapted to study information dynamics in other domains. For instance, one could explore how the relationship between movie buzz and performance equates with that of other goods and services such as video games, restaurants, musical artists and tourist attractions. An even more interesting exploration would be the potential application to popularity of extremist groups, both terrorist and legitimate political groups that are far from the median voter. For this latter case, a study might draw an analogy between a population of homogenous films and a collection of homogenous extremist groups, with information about each group being disseminated among a network of disenfranchised individuals and the 'winning' groups being the ones that can garner the greatest buzz. Such a study would first test the hypothesis that the analogy holds, and if so, then one might test what happens if changes are made to the structure of the network to disrupt the flow of information.

——••••——

[Carl Spackler]: *Licensed to kill gophers by the government of the United Nations. A man, free to kill gophers at will. To kill, you must know your enemy, and in this case my enemy is a varmint. And a varmint will never quit - ever. They're like the Viet Cong - Varmint Cong. So you have to fall back on superior intelligence and superior firepower. And that's all she wrote.*

Quote from the movie *Caddyshack*

REFERENCES

Albert, S. (1998). Movie Stars and the Distribution of Financially Successful Films in the Motion Picture Industry. *Journal of Cultural Economics*, *22*, 249–270.

Allsop, D., Bassett, B., & Hoskins, J. (2007). Word-of-Mouth Research: Principles and Applications. *Journal of Advertising Research*, *December 2007*.

Austin, B. A. (1989). *Immediate Seating: A Look at Movie Audiences*. Belmont, CA: Wadsworth.

Barabási, A.-L., & Albert, R. (1999). Emergence of Scaling in Random Networks. *Science*, *286*, 509–512.

Basuroy, S., Chatterjee, S., & Ravid, S. A. (2003). How Critical Are Critical Reviews? The Box Office Effects of Film Critics, Star Power, and Budgets. *The Journal of Marketing*, *67*, 103–117.

Bhattacharya, P., & Chakrabarti, B. K. (2009). Fractal Models of Earthquake Dynamics. In *Reviews of Nonlinear Dynamics and Complexity*. Germany: Wiley-VCH.

Bianconi, G., & Barabási, A.-L. (2001). Bose-Einstein Condensation in Complex Networks. *Physical Review Letters*, *86*, 5632–5635.

Box Office Mojo. (2014). Retrieved from http://www.boxofficemojo.com/

Burzynski, M., & Baker, D. (1977). The Effect of Positive and Negative Prior Information on Motion Picture Appreciation. *Journal of Social Psychology*, *101*, 215–218.

Cabral, L., & Natividad, G. (2014). *Box Office Demand: The Importance of Being #1*. New York University.

Canterbery, E. R., & Marvasti, A. (2001). The U.S. Motion Pictures Industry: An Empirical Approach. *Review of Industrial Organization*, *19*, 81–98.

Cilliers, P. (1998). *Complexity and Postmodernism: Understanding Complex Systems*. London, England: Routledge.

Collins, A., Hand, C., & Snell, M. C. (2002). What Makes a Blockbuster? Economic Analysis of Film Success in the United Kingdom. *Managerial and Decision Economics*, *23*, 343–354.

Dale, E. (1937). Analyzing the Movie Market. *Educational Research Bulletin*, *16*, 212–216.

De Vany, A. (2004). *Hollywood Economics: How Extreme Uncertainty Shapes the Film Industry*. London, England: Routledge.

De Vany, A., & Walls, D. (1996). Bose-Einstein Dynamics and Adaptive Contracting in the Motion Picture Industry. *The Economic Journal*, *106*, 1493–1514.

Elberse, A., & Eliashberg, J. (2003). Demand and Supply Dynamics for Sequentially Released Products in International Markets: The Case of Motion Pictures. *Marketing Science*, *22*, 329–354.

Eliashberg, J., Jonker, J.-J., Sawhney, M. S., & Wierenga, B. (2000). MOVIEMOD: An Implementable Decision-Support System for Prerelease Market Evaluation of Motion Pictures. *Marketing Science*, *19*, 226–243.

Eliashberg, J., & Shugan, S. M. (1997). Film Critics: Influencers or Predictors? *The Journal of Marketing*, *61*, 68–78.

Gemser, G., Van Oostrum, M., & Leenders, M. (2007). The Impact of Film Reviews on the Box Office Performance of Art House Versus Mainstream Motion Pictures. *Journal of Cultural Economics*, *31*, 43–63.

Goldman, W. (1983). *Adventures in the Film Trade*. New York: Warner Books.

Hahn, M., Park, S., Krishnamurthi, L., & Zoltners, A. A. (1994). Analysis of New Product Diffusion Using a Four-Segment Trial-Repeat Model. *Marketing Science*, *13*, 224–247.

Hennig-Thurau, T., Houston, M., & Sridhar, S. (2006). Can Good Marketing Carry a Bad Product? Evidence From the Motion Picture Industry. *Marketing Letters*, *17*, 205–219.

Hennig-Thurau, T., Houston, M., & Walsh, G. (2007). Determinants of Motion Picture Box Office and Profitability: An Interrelationship Approach. *Review of Managerial Science*, *1*, 65–92.

Herrera, C., & Zufiria, P. J. (2011). Generating Scale-Free Networks with Adjustable Clustering Coefficient Via Random Walks. *arXiv Preprint arXiv:1105.3347*.

Jedidi, K., Krider, R. E., & Weinberg, C. B. (1998). Clustering at the Movies. *Marketing Letters*, *9*, 393–405.

Jones, J. M., & Ritz, C. J. (1991). Incorporating Distribution into New Product Diffusion Models. *International Journal of Research Marketing*, *8*, 91–112.

Karniouchina, E. V. (2011). Impact of Star and Movie Buzz on Motion Picture Distribution and Box Office Revenue. *International Journal of Research Marketing*, *28*, 62–74.

King, T. (2007). Does Film Criticism Affect Box Office Earnings? Evidence From Movies Released in the U.S. in 2003. *Journal of Cultural Economics*, *31*, 171–186.

Lehmann, D. R., & Weinberg, C. B. (2000). Sales through Sequential Distribution Channels: An Application to Movies and Videos. *The Journal of Marketing*, *64*, 18–33.

Levin, A. M., & Levin, I. P. (1997). Movie Stars and Authors as Brand Names: Measuring Brand Equity in Experiential Producs. *Advances in Consumer Research*, *24*, 175–181.

Lilien, G., Rao, A. G., & Kalish, S. (1981). Bayesian Estimation and Control of Detailing Effort in a Repeat Purchase Diffusion Environment. *Management Science*, *27*, 493–506.

Liran, E. (2007). Seasonality in the U.S. Motion Picture Industry. *The Rand Journal of Economics*, *38*, 127.

List of highest paid film actors. (2014). Retrieved November 25, 2014, from http://en.wikipedia.org/w/index.php?title=List_of_highest_paid_film_actors&oldid=627417661

Litman, B. R. (1982). Decision Making in the Film Industry: The Influence of the TV Market. *Journal of Communications*, *32*, 33–52.

Litman, B. R. (1983). Predicting Success of Theatrical Movies: An Empirical Study. *Journal of Popular Culture*, *16*, 159–175.

Litman, B. R., & Ahn, H. (1998). Predicting Financial Success of Motion Pictures. In B. R. Litman (Ed.), *The Motion Picture Mega-Industry*. Needham Heights, MA: Allyn & Bacon.

Litman, B. R., & Kohl, L. S. (1989). Predicting Financial Success of Motion Pictures: The '80s Experience. *Journal of Media Economics*, *2*, 35–50.

Liu, Y. (2006). Word of Mouth for Movies: Its Dynamics and Impact on Box Office Revenue. *Journal of Marketing*, *70*, 74–89.

Mislove, A., Marcon, M., Gummadi, K. P., Druschel, P., & Bhattacharjee, B. (2007). Measurement and Analysis of Online Social Networks. In *IMC'07*. San Diego, CA.

Mizerski, R. W. (1982). Viewer Miscomprehension Findings Are Measurement Bound. *Journal of Marketing*, *46*, 32–34.

Moon, S., Bergey, P., & Iacobucci, D. (2009). Dynamic Effects of Movie Ratings on Movie Revenues and Viewer Satisfaction.

Moul, C. C. (2001). *Word-of-mouth and saturation: Why movie demands evolve the way they do.* Working Paper, Department of Economics, Washington University, St Louis MO.

MPAA. (2013). *Motion Picture Association of America Theatrical Market Statistics*.

Muser. (2011). *Turning the Silver Screen to Gold: An Analysis of Opening Weekend Box Office Success*. Eastern Illinois University, Charleston, IL.

Naylor, T. H., & Finger, J. M. (1967). Verification of Computer Simulation Models. *Management Science*, *14*, B92–B106.

Neelamegham, R., & Chintagunta, P. (1999). A Bayesian Model to Forecast New Product Performance in Domestic and International Markets. *Marketing Science*, *18*, 115–136.

Neelamegham, R., & Jain, D. (1999). Consumer Choice Process for Experience Goods: An Econometric Model and Analysis. *Journal of Marketing Research*, *36*, 373–386.

Nelson, P. (1970). Information and Consumer Behavior. *Journal of Political Economy*, *78*, 311–329.

NetLogo References. (2016). Retrieved from http://ccl.northwestern.edu/netlogo/references.shtml

Nolan, J. P. (1999). *Fitting Data and Assessing Goodness-of-Fit with Stable Distributions." Applications of Heavy Tailed Distributions in Economics, Engineering and Statistics*. American University, Washington, DC.

Pangarker, N. A., & Smit, E. V. D. M. (2013). The Determinants of Box Office Performance in the Film Industry Revisited. *South African Journal of Business Management*, *44*.

Popper, K. (1959). *The Logic of Scientific Discovery*. New York, NY: Basic Books.

Prag, J., & Casavant, J. (1994). An Empirical Study of the Determinants of Revenues and Marketing Expenditures in the Motion Picture Industry. *Journal of Cultural Economics*, *18*, 217–235.

Proykova, A., & Stauffer, D. (2002). *Social Percolation and the Influence of Mass Media*. Sofia-1126, Bulgaria, D-50923 Koln, Euroland: Department of Atomic Physics, University of Sofia. Retrieved from http://arxiv.org/pdf/cond-mat/0203375.pdf

Ravid, S. A. (1999). Information, Blockbusters, and Stars: A Study of the Film Industry. *The Journal of Business*, *72*, 463–492.

Sargent, R. G. (2011). Verification and Validation of Simulation Models. In *Proceedings of the 2011 Winter Simulation Conference* (pp. 183–198). Phoenix, AZ.

Sawhney, M. S., & Eliashberg, J. (1996). A Parsimonious Model for Forecasting Gross Box-Office Revenues of Motion Pictures. *Marketing Science*, *15*, 113–131.

Schechtman, O. (2013). The Coefficient of Variation as an Index of Measurement Reliability. In *Methods of Clinical Epidemiology*. Berlin, Germany: Springer.

Schmidt, D., Best, J., & Blumberg, M. S. (2011). Random Graph and Stochastic Process Contributions to Network Dynamics. *American Institute of Mathematical Sciences*, *2011*, 1279–1288.

Sinha, S., & Raghavendra, S. (2004). Hollywood Blockbusters and Long-Tailed Distributions. *The European Physical Journal B - Condensed Matter and Complex Systems*, *42*, 293–296.

Sochay, S. (1994). Predicting the Performance of Motion Pictures. *Journal of Media Economics*, *7*, 1–20.

Wallace, W. T., Seigerman, A., & Holbrook, M. B. (1993). The Role of Actors and Actresses in the Success of Films: How Much Is a Movie Star Worth? *Journal of Cultural Economics*, *17*, 1–27.

Weekly Box Office March 18-24 2016. (2016). Retrieved April 1, 2016, from http://www.boxofficemojo.com/weekly/chart/?yr=2016&wk=12&p=.htm

What's the Tomatometer™? (2015). Retrieved July 6, 2015, from http://www.rottentomatoes.com/

Wilensky, U. (1999). *NetLogo*. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University. Retrieved from http://ccl.northwestern.edu/netlogo/

Yearly Domestic Gross. (2016). Retrieved from http://pro.boxoffice.com/statistics/yearly

Zufryden, F. S. (2000). New Film Website Promotion and Box-Office Performance. *Journal of Advertising Research*, *40*, 55–64.

BIOGRAPHY

Holly Russo served in the US Air Force for six years after graduating from Absegami High School in 1985. Upon receiving her honorable discharge and living in England for a few years, she returned to settle in Northern Virginia and began working in research, development and analysis for various organizations including RAND, Center for Naval Analyses, BAE Systems, Joint Warfare Analysis Center, i_SW Corporation and currently Deloitte & Touche LLP. She received her Bachelor of Science in Economics from George Mason University in 2002, and her Master of Arts in Economics also from GMU in 2005. She also holds the INFORMS Certified Analytics Professional certification, attained in 2015. She can be contacted via LinkedIn profile https://www.linkedin.com/in/hollyarusso.