TNO Information and Communication Technology

Understanding Consumer Behaviour through Data Analysis and Simulation

Are social networks changing the world economy?

Master's Thesis

Date	August 10, 2008	
Author	Jeroen Latour	
Supervisors	Betsy van Dijk (chair)	University of Twente
	Mannes Poel	University of Twente
	Dirk Heylen	University of Twente
	David Langley	TNO ICT
	Wander Jager	University of Groningen

Number of pages 72

Contents

1 Introduction

This thesis will analyse the music choices of users of Last.fm, a social networking site focused on music. Their choices will be the basis for an exploration of the role of social relationships in music consumers' choice of new songs to listen to. Agent-based modelling, a form of simulation modelling, will be the primary tool used here. The study should serve as an example of how detailed data such as the data on music choice – which is increasingly becoming available – can enable new agent-based modelling research and lead to new insights into such social phenomena as consumer behaviour.

This chapter will describe the field of agent-based modelling, show the potential of using detailed data sets in this type of research, introduce the research questions examined in this thesis and give an overview of the structure of this document.

1.1 Overview

An agent-based computational model allows researchers to simulate the outcome of complex interactions. Simulations are done by creating a virtual environment in which a large number of autonomous agents operate. Each of the agents follows a micro-level specification that is relatively simple, but when brought together they can interact in highly complex ways (Epstein, 1999).

This modelling methodology has a long history, dating back to the work of John von Neumann with his "universal constructors" and "cellular automata" (von Neumann & Burks, 1966). These small programs could interact and reproduce and as such were capable of forming a virtual society... at least in theory. Even though authors like Thomas Schelling already suggested using these automata to model the social sciences (Schelling, 1978), at that time the computing capacity was too limited to put these ideas in practice (Epstein & Axtell, 1996).

In the last few decades, advances in computing power have caused a surge of interest in agent-based models (ABMs). Increasingly, the methodology has been used to work on the many challenges of the social science. In particular, Arthur (1991) and Holland & Miller (1991) introduced the technique to economic modelling. Agent-based models, they argued, would not just model the virtual society's behaviour at a micro-level, but also attempt to uncover the motivations and processes underlying that behaviour.

Such a model could be more flexible in dealing with a changing environment, having these motivations to fall back on. As such, it could be resistant against what has become known as the Lucas Critique. As Lucas (1976) pointed out in his policy evaluation, it is doubtful that any models based on high-level aggregate data will remain valid. His reasoning was that such models would be unable to accurately incorporate changes in the environment. He criticized the use of macro-level models as a basis for policy change, as they would alter the 'rules of the game'.

Since ABMs often specify behaviour at a higher level, their specification may not be dependent on these rules of the game, and they might be able to more accurately predict how agents respond in the new situation. However, as Garcia (2005) and others have noted, ABMs have limited value in predicting the future. Instead, they are often used as tools to understand the dynamics of a society and to help refine existing theory (Bonabeau, 2002).

ABMs are usually trained and validated on macro-level data, which describes the behaviour at an aggregate level. However, this approach often requires many assumptions and it is impossible to show or even to make a credible case that a model is an accurate representation of reality. Even though a model might generate the appropriate behaviour at a macro-level, there could be several micro-level specifications that do so, not all of which are correct (Epstein, 1999).

Proving correctness of an ABM seems impossible without having complete insight into the system that is being modelled. In many domains, this insight cannot be reached, because the agents' motivations are not transparent. This is especially true with modelling consumer behaviour. Even if consumers would try to truthfully explain the reasons behind their actions, they might not be able to tell the whole story. Decisions are influenced not only by external factors, but also by personal desires and goals, often at a subconscious level. Consumers have to guess about their motivations as much as or even more than an outside researcher.

Increasingly, there is a new type of data available that could benefit researchers training, validating and applying agent-based models. This type of data, micro-level data, describes behaviour at an individual level, as opposed to the aggregate statistics commonly used up to this point. As such, it provides behavioural information on how individuals react to various situations. By describing these situations, a researcher can trace back which information was available at the time of the action, and determine which features affect the decision taken. This data is becoming available because companies are keeping extensive records on their customers and because consumers themselves increasingly choose to share vast amounts of information through social networks.

This thesis will explore one such micro-level data set, on the music choices of a group of Last.fm users. This data set describes in detail which songs each user chose to listen to, at what time they first listened to the song, which of their friends had already listened to the song and how well the song matches with what they usually listen to. This data will be used to examine the processes underlying the adoption of new songs by music consumers. Its results will show the value of this type of data, and illustrate the type of analysis that can be performed on it.

1.2 Research questions

This study will explore the value of agent-based modelling in combination with micro-level data sets. It will do so through applying this technique on this type of data to examine a popular topic in marketing research: social contagion. The theory of social contagion suggests that people's adoption of new products is a function of their exposure to other people's knowledge, attitudes or behaviours concerning the new product. Researchers have introduced several theoretical accounts of social contagion, including social learning under uncertainty, social-normative pressures, competitive concerns and performance network effects (van den Bulte & Stremersch, 2004).

This multitude of explanations and studies suggests this area is still very much being researched. Many have performed studies examining social contagion in a wide range of domains and markets. Combined, these studies enable such meta-analytic studies as that of van den Bulte & Stremersch (2004), who compared the result of 54 publications to determine which of the conclusions reported by these publications were robust over multiple studies and markets. The success of these meta-analytic studies shows that the value of additional studies examining social contagion in a specific domain extends beyond that domain.

Contributing to these efforts, this thesis will explore the role of social contagion in consumers' adoption of new music. It will do so by exploring the diffusion of new songs among users of the Last.fm network. This social network tracks which songs its members listen to and publishes this information on their online 'profiles'. In combination with information on relationships between these users, this results in a socially connected micro-level data set that could be very valuable in adoption research. The role of social contagion in the adoption of new songs will be the first question examined in this thesis.

Closely related to the study of social contagion is the study of heterogeneity in consumer influence. Over the years, many researchers have attempted to identify groups of consumers with a unique role in the diffusion of a new product. The most famous example of this is the categorization of adopters as (1) early adopters, (2) early majority, (3) late majority or (4) laggards (Ryan & Gross, 1943), based on the time of adoption relative to all other adopters. Valente (1996) later revisited this concept and redefined the categories locally, using the time of adoption relative only to the social circle.

These categorizations lead to the question why some people adopt much sooner than others. This question has inspired the influentials theory (van den Bulte & Wuyts, 2007), which suggests that consumers can be divided into *influentials* and *imitators*. Influentials are more in touch with new developments, and their behaviour is thought to influence the group of imitators. This theory has until now not been unanimously accepted and there is a need for studies examining the existence and characteristics of a group of influentials who influence the adoption by imitators. This thesis will explore whether groups of influentials or imitators could be found, and attempt to discover their characteristics. This will be the second question explored here.

Assuming that social contagion does play a role in the adoption of new songs, and some users are influenced by their friends choice of music, a final question remains. What exactly is the effect of social contagion on the diffusion of a new song? How would this effect change if social contagion increases or decreases over time? This will be the third and final question examined in this thesis.

To summarize, this thesis will explore the following questions:

1. What is the role of social contagion in the adoption of new songs by users of the Last.fm network?
2. Who are the influencers and the imitators on the Last.fm network?
3. How does social contagion change the adoption of new songs by users of the Last.fm network?

1.3 Document structure

This chapter has explored the field of agent-based modelling research, and listed current issues with both training, validating and applying agent-based models. It has introduced the concept of micro-level data – data that describes behaviour at an individual level – and has put forward that this type of data is becoming increasingly available. As a result, it has set a task of exploring how this type of data could benefit agent-based modelling research, by showing what can be gathered from one such data set. This data set, which describes in detail which songs a group of Last.fm users have listened to, will be used to explore the role of social contagion in the adoption of new songs.

The remainder of this thesis will attempt to answer three research questions identified in the previous section. First, chapter 2 will describe how this data set will be used to build and validate an agent-based model, and how this model and the data will be used to answer the research questions. The next chapters will describe the preparations that were necessary to get to examining the role of social contagion. In turn, they will describe the data (chapter 3), model (chapter 4) and the validation of model and data (chapter 5).

Chapter 6 will use the model and data set to answer the research questions. Finally, chapter 7 will present conclusions on the role of social contagion in the adoption of music, while chapter 8 present directions for future research and discuss what this study has shown about the merit of using agent-based models and micro-level data.

2 Method

This chapter will describe the approach taken in this study. First, it will describe the approach to build an agent-based model for the adoption of new songs by Last.fm users. Then, it will describe how this model and the data set are used to answer the research questions.

2.1 Building an agent-based model

The method for building an agent-based model can be divided into three steps:

1. Collecting the data
2. Calibrating the agent-based model
3. Validating the agent-based model

This section will describe the approach taken to each of these steps.

2.1.1 Data

Micro-level data can be very useful in agent-based modelling research. However, such data is usually not readily available and will need to be collected first. To do so, three things need to be done:

1. Collect observations of ’actions’
2. Collect the state of the environment at the time of each action
3. Collect information on the network structure

For the purpose of studying music diffusion, the ’actions’ observed in the first step can be defined as someone listening to a song for the first time. Last.fm keeps track of all the songs its users listen to and publishes these records on the users’ profiles. This makes ’collecting observations’ as simple as going through these records to find songs that the user had never played before.

The second step is to collect the state of the environment at the time of each action. This means collecting all the information the user had at their disposal when they decided to listen to a song. Obviously, ’all information’ is quite a broad definition and implies collecting a sheer unlimited amount of data, but we need to concern ourselves with only a small portion of it. Since the purpose of this data is to help calibrate an agent-based model, only the information used by the model to make a decision is required.

This study will use a model proposed by Delre (2007). This model will be described in detail in chapter 4. For now, it will suffice to say that agents in this model base their decision on two inputs: one product-related, one social.

1. The product-related input measures the match between the song and the user's taste. This information will be based on a measure of similarity between this song and the songs most commonly played by the user, provided by Last.fm algorithms.

2. The social input measures how many of a user's friends have already listened to the song. This information will be collected by retrieving a list of the user's Last.fm friends and analysis their listening records.

These lists of Last.fm friends are also needed for the third step, collecting information on the network structure. Agent-based models simulate the behaviour of a network of actors[1] and past research has shown how much the network structure can affect the simulation outcome (Bonabeau, 2002). Because of this, it is important to record who is friends with who, to facilitate structural analysis in the calibration stage.

This section has described in overview how the data set will be collected: by recording occurences of the first time people listen to a song, recording what the state of their environment was at that time and collecting information on the network structure. Chapter 3 will describe the data collection process in more detail.

2.1.2 Calibration

Several studies have presented step-by-step approaches to developing an agent-based model. Garcia (2005) proposed a fairly complete approach, identifying the following steps:

1. *Theory Operationalization and Cognitive Map Creation.* Determining which of the system's elements are important enough to be included in the model.

2. *Agent Specification.* Defining the strategies and characteristics governing an agent's behaviour.

3. *Environmental Specification.* Defining the structure and behaviour of the environment.

4. *Rules Specification.* Defining the rules of the game.

[1]Granted, some agent-based models simulate a two or three-dimensional world, but this could be viewed as a network where connections are defined by spatial proximity.

5. *Measurements Recording.* Defining which results should be recorded as the model's output. Since ABMs are often used to examine aggregate or emergent behaviour, these results are usually measured globally.

6. *Run Time Specification.* Defining how many iterations will be simulated and how many runs are needed to get stable results.

This study will use a pre-existing model for the diffusion of innovation, proposed by Delre (2007, see chapter 4). His specification largely coveres steps 1 through 4. However, Delre specified the agents and the environment in a parameterised manner. Parameterisation creates a distinction between agent attributes and behaviour (Twomey & Cadman, 2002). This approach allows for agent behaviour to be specified in general terms, with the attributes as 'dials' to finetune behaviour.

Finding the appropriate parameters will be done in two parts:

1. *Agent Specification.* Finding the agent parameters that best fit the observed behaviour.

2. *Environmental Specification.* Characterizing the network structure so that it can be accurately replicated.

These parts, and the final steps of model construction, will be discussed in more detail below.

Agent specification

The micro-level data set could be considered as a set of 'cases', each describing a decision taken by one of the agents (in this case, whether or not to listen to a song). By considering the circumstances under which an agent takes each of the possible decisions, it is possible to determine the behaviour parameters that best approximate the behaviour seen in the data set.

This regression task is very common in machine learning literature (Alpaydin, 2004) and as such, the steps proposed here are quite similar to the steps taken to approach a classification problem in machine learning.

Determining agent characteristics from a set of real-world cases

1. Extract cases from the data set
2. Group cases by user and decision
3. For each decision, determine the state of the environment at the time
4. Determine the parameters that best predict the decision, based on the state of the environment

5. Analyse the distribution of parameter values

It seems valuable to take a moment to apply these steps to a simple example. Consider a simple agent-based model, in which each agent focuses only on preventing loneliness. Every time step, it looks around to see how many other agents are in the room. If the number of agents in the room is at least as high as the agent's loneliness threshold, the agent is happy and stays to chat with the other agents. If the number of agents drops below that level, the agent gets lonely and moves to another room in search of a bigger crowd.

Suppose a micro-level dataset for this model tells us that:

- When 4 agents were in the room, agent Alice moved to another room.
- When 3 agents were in the room, agent Bob moved to another room.
- When 6 agents were in the room, agent Alice happily stayed to chat with the other agents.
- When 2 agents were in the room, agent Bob moved to another room.

Each of these observations could be interpreted as cases, describing which decisions were made by the agents in which conditions (step 1). They can be grouped by user and decision (step 2), resulting in 1 'move' case and 1 'stay' case for Alice, and 2 'move' cases for Bob. For each of these cases, we already know the state of the environment (step 3), namely the number of agents that were in the room when Alice or Bob made their decision.

For each of the agents, we can try different parameter values and see how many times the model will predict the same decision as we saw in the cases (step 4). If we choose a loneliness threshold of 4 for Alice, she would stay in the room both when there were 4 agents and when there were 6 agents in the room. Since the data set shows that Alice left with 4 agents in the room, the agent correctly predicted 50% of the decisions. With a threshold of 7, she would always leave the room, which does not match the observation of her staying with 6 agents (again, 50% precision). With a threshold of 5 or 6, all the decisions are correctly predicted. The data set cannot tell us whether the threshold is 5 or 6 so until more data becomes available, either option is equally likely.

Bob has even less data available to help us determine his loneliness threshold. Since we only have 'move' cases for him, we could set the loneliness threshold to any number higher than 3 and achieve 100% precision. This introduces a very large margin of error, and illustrates the importance of excluding these users from the case set.

Of course, models are a simplication of reality and real-world data will give conflicting signals about the 'true' value of the loneliness threshold. For this reason, it will rarely be possible to find values that give 100% precision and an agent-based modeller will have to be satisfied with a best-possible fit. The quality of the fit will be one of the things to consider when analysing the validity of the model.

In this case, applying the steps only suggests parameter values for one of the two agents. If more parameter sets were found, the next step (step 5) could be to find a way to aggregate these parameter values, for example by finding a statistical distribution whose samples are similar to the parameter values.

The five steps presented here illustrate how a regression algorithm can be applied to a micro-level data set. Much the same as was done for the example data set, a regression algorithm can find the optimal set of parameters to fit the model to the data set.

Environmental specification

The second part of the parametrisation is the environmental specification, in particular determining the structure of the network of agents. Some of the properties of networks can be represented by simple mathematical models that interpolate between a completely structured and a completely random graph. A completely structured graph is a regular lattice, where each node is connected to its k nearest neighbours on the lattice. The randomness r is defined as the fraction of links in the lattice that were randomly rewired to get the network structure that is being described (Watts & Strogatz, 1998).

Real-world networks are neither completely ordered nor completely random. This hybrid type of network, with properties of both completely random and completely ordered networks, is often referred to as a 'small-world network' (Barabási & Albert, 1999). Amaral et al. (2000) present three classes of small-world networks, characterized by the degree distribution (the number of people with 0 friends, the number with 1 friend, and so on):

1. Scale-free networks, characterized by a degree distribution with a tail that decays as a power law. In this case, the tail of the distribution would fall on a straight line in a log-log plot.

2. Broad-scale networks, characterized by a degree distribution that has a power law regime followed by a sharp cut-off, like an exponential or Gaussian decay of the tail. In this case, the distribution would initially show a straight line in a log-log plot, but the tail decays faster than this power law.

3. Single-scale networks, characterized by a degree distribution that does *not* have a power law regime, but does have a fast decaying tail (exponential, Gaussian).

Each of these classes have been studied to determine which mechanisms cause such a structure. For example, for scale-free networks Barabási & Albert (1999) showed that this behaviour is necessarily the consequence of two generic mechanisms:

1. Networks expand continuously with the addition of new vertices.
2. New vertices attach preferentially to sites that are already well-connected.

This shows that determining the type of network can also provide information about the dynamics of the network. With that, its value extends beyond creating a realistic structure of agents. By analysing the necessary conditions for a certain structure to emerge, as Barabási & Albert have done, we gain insight into how the network is shaped. For example, if a social network site was found to have a scale-free network, it would show that new users are more likely to connect to active users in the community.

Calculating the degree distribution is the most obvious route to determining the type of network, as all of the types identified above have degree distributions with distinct characteristics. Once the type of network has been determined, network structure can be further determined by determining the exponents of the distribution, such as the exponent of the power law in the case of a scale-free network.

Determining network structure in a socially connected data set

1. Calculate the degree distribution of all individuals in the data set.
2. Show this distribution in a log-log plot to visually determine the type of network, according to the characteristics listed above.
3. Use curve fitting algorithms to determine the model parameters of either a power law or an exponential distribution, depending on the network class.
4. Use existing algorithms to create the desired network class with the appropriate parameters, recreating the degree distribution.

Final Steps

The work done by Delre (2007), in combination with the parameterisation described in the previous steps, covered steps 1 through 4 of the procedure proposed by Garcia (2005) and described at the start of this section. Two steps remain: measurements recording and run time specification.

In the measurements recording step, the researcher specifies what in the simulation will be recorded. Ideally, every action taken during the simulation will be recorded and compared to the actions recorded in real life. However, this is not possible in most cases, as the network of agents will be created from a generalized description of the network structure. As such, there is no one-to-one relationship between simulated agents and real-life people.

Alternatively, simulation output will be recorded at an aggregate level. The simulation will record adoption at the end of every day (that is, once every timestep), which can then be compared to real-life adoption at the end of that day. This will provide more detailed statistics than simply recording adoption at the end of the simulation, as this will also enable studying how adoption develops during the simulation.

A proper run time specification ensures that simulation results are unlikely to be influenced by chance. If only one simulation run was being done, an unlucky choice of initial adopter could for example stop adoption at day 1, while almost every other time the song will go on to become a major hit. In this study, at least ten simulations are done for every product that is being tested, and the average adoption is taken at every time step. If at the end of those ten simulations, the averages have not stabilized (defined as: the integer values of the average changed after the last simulation), simulations continue until a maximum of fifty runs for that product. This range of run counts was used to cut down simulation time for those products that have very stable simulation outcomes.

With that, the last of the six steps identified by Garcia (2005) has been covered. This subsection has provided an overview of how a parameter set and a description of the network structure are extracted from the data, and combined with the existing model specification. This process will be described in more detail in chapter 4.

2.1.3 Validation

This subsection will describe the validation of the model constructed using the method described in the previous subsection. Three major approaches to validation are usually distinguished (Carley, 1996; Fagiolo et al., 2005; Windrum et al., 2007; Garcia et al., 2007).

1. The indirect calibration approach (Dosi et al., 2006) first performs validation to find ranges of parameter values that produce 'valid' output. Only after validation has been completed is the model calibrated (hence the indirectness of the approach), choosing the parameter values from the ranges determined in the first step.

2. The Werker-Brenner approach (Werker & Brenner, 2004) starts with calibration: using existing empirical knowledge to calibrate initial conditions and the ranges of model parameters. Step two is to validate this calibrated model to further reduce the parameter space. The final step involves a further round of calibration, with help from historians.

3. The history-friendly approach (Malerba et al., 1999; Malerba & Orsenigo, 2001) does not focus on defining the parameter space that leads to valid output, but rather on finding the parameter values that best fit the observed data. As such, it is quite different from the other two. This method attempts to find the parameters that reproduce the output seen in the data set, as measured by a number of 'stylised facts'. These are simply the 'measurements' mentioned in Garcia (2005)'s method for building agent-based models.

From these approaches, the history-friendly approach is most appropriate for finding the parameter values that generate the closest approximation of observed reality, as is testified by the descriptions. It can be applied using aggregate adoption figures to measure the quality of the fit – the 'stylised facts' referred to above – to find the optimal solution. Obviously, it is highly unlikely that any combination of parameter values will be able to reproduce all user's listening choices perfectly, as the history-friendly approach seems to suggest. Instead, this process will be limited to finding the best possible match.

As becomes apparant from the descriptions of each of these approaches, calibration and validation go hand in hand. The history-friendly approach can be used to perform calibration as well. In fact, it prescribes running simulations for every combination of parameters to find the optimal solution. As this is too computationally intensive to be feasible, a two-step approach will be used here. Calibration will be performed as described in the previous section, using the fit to the training data to produce a candidate parameter set. Several candidates can be produced to examine the effects of changing assumptions. For example, there could be several candidates to try different approaches to describe the network structure. The final choice among these candidates is done in the validation step.

Validation, using the history-friendly approach

1. Implement the agent-based model
2. Initialise agents and environment
3. Simulate adoption of all products in the test set
4. Repeat steps 2 & 3 several times to filter random deviations
5. Compare simulated adoption figures to real data

Validation is the last step in building an agent-based model. Collecting the data and calibrating and validating the model will have produced an agent-based model that can be used to examine the role of social contagion in music adoption.

2.2 Examining social contagion

Chapter 1 identified three research questions:

1. What is the role of social contagion in the adoption of new songs by users of the Last.fm network?
2. Who are the influencers and the imitators on the Last.fm network?
3. How does social contagion change the adoption of new songs by users of the Last.fm network?

This section will describe how each of these questions will be explored.

2.2.1 Role of social contagion

The current role of social contagion in the adoption of new songs by users of the Last.fm network will be explored in two ways: through exploring the results of model calibration and through exploring the q/p ratio.

The results of the model calibration can give us valuable insights about the role of social contagion. In particular, a lot can be learned from the balance between social influences and product-related influences that best recreates the data set. If the balance is mostly towards product-related influences, this suggests that the adoption by friends was not an accurate predictor of adoption. This, in turn, suggests that social contagion does not play a big role for that user.

There is a second way to measure social contagion, one that can also be used in cases where product-related influences are so strong that any signs of social influence all but disappear. This alternate method builds on the body of work following Bass (1969). Bass proposed an aggregate product growth model where the predicted number of adopters in any timestep is based on the number of adopters in the previous timestep. His work has been used by many authors.

In particular, many studies have focused on the q/p ratio, referring to the q and p in Bass' equation for the rate of adoption (see chapter 4). In this equation, q was a measure of social influence and p was a measure of product influences. Thus, the q/p ratio provides a measure of the importance of social influence in a given market, and some have used it to compare the social effect in various domains (van den Bulte & Stremersch, 2004).

With so many studies examining the q/p ratio in their data set, fitting the data to the Bass equations and calculating the q/p ratio allows for comparison of the social influence in the researcher's data with previous studies (van den Bulte & Stremersch, 2004). The general approach to calculating the q/p ratio is to fit the rate of adoption to Bass' equation ($r(t) = p + qF(t)$).

However, this is only applicable when analysing a single diffusion curve. There is no commonly used definition for an aggregate q/p ratio, for a collection of diffusion curves. Jager (2008)'s definition is used here: q/p equals the number of adoptions with social influence divided by the number of adoptions without social influence. Here, an adoption with social influence is an adoption that was preceded by an adoption of a friend or connection. Low values for this ratio would indicate social contagion plays a small role in music adoption, while high values indicate social contagion plays a big role.

2.2.2 Influentials and Imitators

If social contagion plays a role in music adoption, it is not improbable that some people will be more easily influenced by their friends. Similarly, it seems quite possible that some people will have more of their friends following them.

To determine whether there is a group of users who has more influence than others, one must first define 'influence'. Here, influence is based on the number of friends who adopt later. The expected number of friends who adopt later is the number of friends who have not yet adopted multiplied by the probability that a random Last.fm user adopts this song. Influence is then defined as the actual number of friends who adopt later, divided by the expected number. With this definition, influence is compensated for the number of friends. If this were not the case, we would always expect people with many friends to have high influence, and any relation between the two could be discounted for that reason.

This research question will be answered by calculating the influence for all observed adoptions and determining whether any group of users has above-average influence (greater than one). In particular, it is of interest to examine whether people with above-average influence ('influentials') have more friends than other users.

Similarly, it is possible to examine the characteristics of 'imitators', users for whom friends play an important role in determining their music choice. This group can be discovered by analysing the result of model calibration. In particular, the parameter determining the relative importance of social inputs and product-related inputs can signal the degree to which a user is affected by his social circle. Determining which users are influenced by their friends could help find common characteristics and understanding what type of users responds to social contagion in their choice of music.

2.2.3 The effect of social contagion

Assuming that social contagion does play a role in the adoption of new songs, and some users are influenced by their friends choice of music, a final question remains.

What exactly is the effect of social contagion on the diffusion of a new song? How would this effect change if social contagion increases or decreases over time?

This effect will be examined with the agent-based model that was calibrated and validated using the methods described in the first half of this chapter. The simulations will be rerun several times with different model parameters, varying the number of users who are influenced by social contagion. The resulting adoption curves can be compared to determine what effect increased social contagion has on song adoption.

Comparing the outcome of different simulation runs is very similar to comparing the outcome of different simulation runs to real adoption data. The steps presented here are therefor quite similar to the steps of validation:

Investigating the effect of social contagion through simulations

1. Implement the agent-based model
2. Initialise agents and environment
3. Simulate adoption of all products in the test set
4. Repeat steps 2 & 3 several times to filter random deviations
5. Repeat steps 2-4 with different social contagion settings and compare results

There is of course, one important difference. The results of these simulations are only compared to each other, not to real-world adoption data. If the model was validated, it seems fair to assume that the model provides a reasonable representation of consumer behaviour in the music industry. As such, varying the parameters should provide insight into how the market will respond to changing conditions.

This chapter has detailed how data collected from Last.fm will be used to build an agent-based model for the adoption of new songs by Last.fm users. Then, it described how this model and the data set are used to answer each of the three research questions. The following chapters will apply the methods outlined here to collect the data set, build an agent-based model and answer the research questions.

3 Data

This chapter presents a data set describes the listening choices of a set of Last.fm users. The data set is the result of monitoring 34,325 users during 4 months and recording every song that they played. The following sections will introduce Last.fm, describe the collection process and give a sense of the size of the resulting data set.

3.1 Source

Last.fm is a social networking site focused on music. It offers its users software to record which songs they play on their computer, as well as full-length tracks played through the site (but not the radio service offered by Last.fm). This information is transmitted back to the Last.fm servers, where it is used to recommend other music the user might enjoy. The service also generates statistics on the most listened to artists, albums, tracks. These statistics are presented in a profile, which users can share with friends to advertise their taste in music.

The data shared on this website provides amazing insight into people's listening habits. It includes not only a list of artists the user most commonly listens to, but in most cases also a list of the songs listened to in the last few weeks. For each of those songs, it is possible to see when exactly the user listened to that song. In combination with data about social relationships, this enables researchers to see how songs spread through the social network.

Like any social networking site, Last.fm allows users to add these contacts as a 'friend' on their site. This is a bidirectional relationship so it requires mutual consent and both parties will become the friend of the other. After doing that, they are kept informed about the music their friend plays and any music events that he or she is attending. However, this 'friends' feature is not the only way to connect to others through Last.fm.

Each week, their algorithms determine who among their users have a music taste that's closest to what the user listens to. These people (or rather, the top fifty) are presented as 'musical neighbours'. As with friends, Last.fm also informs you about the music these neighbours listen to, helping you to discover new music. From learning new music it is a small step to creating new friends, with Last.fm making it easy to connect to your neighbours and get to know them better. However, it is doubtful that this happens very frequently. As Ellison et al. (2007) explored for

Facebook and others have for different SNSs, online social networks are mainly used to strengthen or maintain offline relationships and very rarely to build new ones.

3.2 Collection

Since its creation in 2002, over 21 million users have signed up for the service (The Guardian, 2008). As all data has to be collected seperately for each user, tracking all users is not possible in the scope of this project. For this reason, the data set was restricted to the music listening habits of a smaller group of users. The data focuses on fans of British independent rock music. This group was chosen because internet plays a big role in independent rock music (CNN, 2006) and its fans are likely to be computer-savvy. For this reason, it seems plausible that selecting this group of Last.fm listeners creates an unbiased selection of fans of independent rock music.

To find these fans, a spider was created to crawl the Last.fm friends network. The spider started from a small seed group of users, randomly selected from the members of the 'Britpop' group[1]. For each of these users, the spider examined all their friends and neighbours, as well as their friends and neighbours. The size of the seed group was chosen such that it could grow no more without making the resulting network too large to revisit biweekly. The revisit was required because listening data was kept for only two weeks before disappearing from the site, and it seemed desirable to have data over a larger period.

For each of the users in the network, the following was collected:

1. A complete profile, including their username, real name, website URL, registration date, age, gender, country, total number of songs played at the time the profile was downloaded and the URLs for their avatar and icon images. This information was not used in the analysis.

2. A full list of their friends and neighbours, at the time the user was first visited.

3. A list of all the songs they listened to, from 18 March 2008 to 18 July 2008. This list was supplemented regularly with new data, as they became available. Every user was examined for new listening data at least every two weeks, to ensure there were no gaps in the data.

4. A list of the user's favorite artists, to help characterise the user's taste. A user's favorite artists are defined as the fifty artists for which the most plays by the user were recorded by the Last.fm service

5. For each of their favorite artists, a list of the artists who are most similar to that artist according to Last.fm's algorithms. These algorithms group those artists that have a largely overlapping fan base (Last.fm, 2008).

[1] http://last.fm/group/Britpop

The following steps were taken to process the data:

1. Songs released before 18 March 2008 were excluded from the data set, because the data on those songs does not include the start of the adoption curve. In those cases, it would have been impossible to determine whether the first recorded instance was really the first time this user listened to that song.

   To determine which songs were released during the observation period, partial listening data for the weeks before March 18 were used. Songs that had more than 2% of its listeners first hearing the song before March 18 were excluded. The filter uses a 2% limit instead of a 0% because it turned out that many newly released songs had one person listening to it a few weeks before release.

2. Users who had listened to less than fifty songs since March 18 were excluded, as these users most likely only let a small portion of their listening behaviour be recorded by Last.fm.

3. All timestamps in the dataset have been converted to the number of days since the first time someone listened to that song (i.e. the release). The model uses discrete timesteps, so this step was necessary to make the progression through time comparable. In the simulations, each timestep will equal one day of simulated time.

3.3 Exploration

In a period of several months, data was collected on 34,325 Last.fm users. Out of the 4.7 million tracks that these users were found to have listened to, over 18,000 had been released between 18 March 2008 and 18 July 2008. These tracks have been listened over 3.4 million times, a million of which for the first time. Over 100,000 adoption cases (instances of a user listening to a song for the first time) were recorded in the seed group (the group of randomly selected users that acted as starting points for the crawler) and their direct friends and neighbours (first degree). Adoptions by friends of friends, friends of neighbours, neighbours of friends et cetera (second degree) have not been included as cases, because the data set did not describe the listening behaviour of all of their friends. The listener behaviour of friends was needed in order to determine how many of a user's friends had listened to the song before the user did.

Last.fm users have a very limited friends list. The users in the extended sample group (see Table 3.1) had a median of 3 and a mean of 9.34 friends. These figures pale compared to the number of friends of a typical Facebook user. In a survey of 4.2 million Facebook users, Golder et al. (2006) found a median of 144 and a mean of 179.53 friends. As in the data of Golder et al., the difference between mean and median is explained by a few users with an exceptionally high number of friends. In the Last.fm data, only 9.9% of the users had more than 21 friends, with one user topping the list at more than 3000 friends.

This large difference in friend counts could partially be accounted by the enormous difference in total network size between Last.fm and Facebook. Facebook reports having over 70 million users (Facebook, 2008) while Last.fm reportedly 'only' has 21 million users worldwide (The Guardian, 2008). With fewer registered users, the odds are lower that users will find their friends on the network.

Still, the difference in size seems too small to completely account for the difference in friend counts. Perhaps the explanation lies with the nature of the service. Since the interaction with friends on Last.fm is mostly limited to music, users might not feel the urge to actively invite all their friends. In fact, perhaps they only talk about music with a small portion of their friends. For all others, they seem more likely to interact with them through Facebook or MySpace.

Table 3.1 gives a more extensive overview of the size of the dataset. Background about the definition of 'cases' in this overview and the reasons for counting these types of cases separately can be found in chapter 4.

6	users in the seed group
3,208	users directly connected to the seed group
3,214	**users in the sample group, for which cases were recorded**
31,111	users connected in the 2^{nd} degree to the seed group
34,325	**users in the extended sample group**
160,255	friend relations (bidirectional)
1,074,402	neighbour relations (unidirectional)
1,234,657	**relations in total**
5,937,770	unique tracks, played by at least one person
18,735	new tracks (released 18/3/2008-18/7/2008, 10+ listeners)
119,758,685	times listened to any track
9,498,011	times listened to new tracks
2,863,556	times listened to new tracks, for the first time
19,094	cases of a user adopting on the first day
3,851	cases of a user adopting the day after a friend did
246,268	cases of a user adopting on a different day
627,311	cases of a user not adopting the day after a friend did
896,524	**cases recorded in total**

Table 3.1: An overview of the dataset

4 Model

This chapter describes the construction of an agent-based model based on the data set presented in the previous chapter. The model will not be constructed from scratch. Instead, an existing model will be used that was built to emulate the diffusion of innovations.

4.1 Background

Throughout the years, an impressive number of models have been proposed to capture the diffusion process of new products and ideas[1]. The most prominent diffusion models were based on the work by Rogers (1976) and Bass (1969). Bass proposed an adoption model based on the assumption that 'the timing of a consumer's initial purchase is related to the number of previous buyers.' With that, his model was the first to incorporate social imitation.

In this model, 'the probability that an initial purchase' – an adoption – 'will be made at time T [...] is a linear function of the number of previous buyers. Thus, $P(T) = p + (q/m)Y(T)$, where p and q/m are constants and $Y(T)$ is the number of previous buyers. Since $Y(0) = 0$, the constant p is the probability of an initial purchase at $T = 0$ and its magnitude reflects the importance of innovators in the social system. [...] The product $q/mY(T)$ reflects the pressures operating on imitators as the number of previous buyers increases.' (Bass, 1969, p. 216) Plotting this function produces the famous 'saddle' curve, with an exponential growth to an adoption peak, followed by an exponential decay.

Bass' model has been extensively tested and performs well in many markets, such as durable goods. Still, many have identified and worked on issues with his model (see Ruiz, 2005, for an overview). One major point is the model's assumption that the group of consumers is homogeneous, which does not hold in every market. Jain et al. (1991) and Hahn et al. (1994) worked on extending the model to an heterogeneous population, but their approach was limited to defining a small number of predetermined groups.

Even with these extensions, macro-level models have to make assumptions about the micro-level behaviour of its consumer population. Most of these models will

[1] An overview of the work on this topic can be found in the surveys by Arts et al. (2006), Mahajan et al. (2000) and Meade & Islam (2006).

assume that the behaviour remains constant over time, or that all consumers behave in the same fashion. The more ambitious who attempt to capture the diversity of the market are soon faced with a model of frightening complexity (Delre, 2007; Chatterjee & Eliashberg, 1990).

Epstein & Axtell (1996) was one of the first to use agent-based models in economics research. He proposed Sugarscape, a resource-allocation model of simple agents 'hiving' virtual sugar mountains. This model showed the potential of using agent-based models to examine market dynamics. Jager (2000) introduced a more generalized model, with agents balancing different needs and resources, and copying successful behaviour from each other. This provided a framework for modelling the basic social interactions, and the consequences of social contagion.

Others have focused not so much on resource allocation, but rather on the diffusion of innovations or new products. In particular, Guardiola et al. (2002) modelled the effects of word-of-mouth marketing, by letting innovators notify their friends of any technology upgrades. Janssen & Jager (2003) combined this word-of-mouth marketing process with a product-oriented version of Jager (2000) and proposed one of the first ABMs simulating market dynamics for new products, not just incremental innovations.

4.2 Model

Building on the work of Janssen & Jager, Delre (2007) proposed several agent-based models for examining innovation diffusion. These models built on the work described in the previous subsection, combining computational models of virus propagation with more realistic models of decision-making and social networks. Each of his agents makes adoption decisions based on a simple weighted utility of individual preference and social influence. The variants are very similar to each other, but differ in minor details to study various dynamics.

In modelling the listening behaviour found in the Last.fm data set (see chapter 3), the author used the model variant Delre described to examine the effects of promotional activities. This model was selected because it simulates not only word-of-mouth promotion, but also mass-media campaigns. Considering the amount of money that is spent on promoting some of the artists, this process seems relevant for an accurate representation of music adoption.

In this model, agents have two thresholds: one for the social influence and one for the product influence. To come to a decision, the agent determines whether enough of his friends have adopted and whether the quality of the product is high enough, both as defined by the thresholds. It then calculates the utility of adopting and proceeds to

Inputs	
a	percentage of adopters in the agent's personal network
q	quality of the product
Parameters	
h	minimum percentage of local adopters required to feel social influence
p	minimum product quality required to feel product influence
β	balance between social and product influence in decision making
U_{min}	minimum utility required for the agent to adopt

Table 4.1: Inputs and outputs of the Delre model

adopt if this utility exceeds the utility threshold. Table 4.1 gives a complete overview of the model's inputs and parameters. The utility function is defined as:

$$U = \beta \cdot x + (1 - \beta) \cdot y$$

$$x = \begin{cases} 1 & \text{if } a \geq h \\ 0 & \text{otherwise} \end{cases}$$

$$y = \begin{cases} 1 & \text{if } q \geq p \\ 0 & \text{otherwise} \end{cases}$$

$$\text{action} = \begin{cases} \textit{adopt and tell all my friends} & \text{if } U \geq U_{min} \\ \textit{do nothing} & \text{otherwise} \end{cases}$$

To perform simulations with this model, the researcher creates a network of agents similar in structure to the real-world situation he is trying to model. An innovation is then seeded, by introducing it to a number of initial adopters, a random sample of e_1 percent of the population. These initial adopters all adopt, regardless of the utility, and tell their friends or connections about the product. The percentage e_1 is defined per product, based on the expected number of initial adopters.

Everyone who hears about the product is given the opportunity to adopt. If they do, they in turn tell all their friends about the product and they all get a chance to adopt in the next timestep. This process models word of mouth promotion. An alternative way of spreading the word about the product is through mass media. Mass-media campaigns can be launched at any time during the lifetime of the product, and give a random sample of e_2 percent of the population a chance to adopt. Contrary to the initial seeding, these new initiates do use their utility function to determine whether to adopt, so they get a chance to ignore the product announcement. If they do adopt, they also tell all their friends, and word of mouth continues. The percentage e_2 is defined per product and per timestep, based on the timing and extent of any promotional campaio

4.3 Exploration

Standard regression algorithms could be used to find the parameter values that produce an optimal fit to the adoptions and non-adoptions for the user. However, these algorithms have great difficulty with the model used here: with three thresholds leading to a binary output, the algorithm has very little information on whether its parameter guess got it closer to finding the optimal solution. With a binary output, other data models such as decision trees could be more suitable for fitting the data, but these would have to replace the model chosen here, abandoning the link to economic theory on consumer behaviour.

A different method will be proposed here. Further inspection of the model revealed that only a small number of values for each of the parameters must be tried to be sure that an optimal solution has been found. This makes the search space small enough to perform an exhaustive search, finding a solution by simply trying all the options and keeping the ones that work best.

For the social threshold h and the quality threshold p, the set of values to try is related to the values of respectively the social influence a and the product quality q found in the data set. These thresholds divide between the values above or equal to a certain point, for which influence is felt, and the values below that point. Of course, the exact cut-point will make a difference when dealing with unseen data, but there is no way to choose between two possible values without analysing more observations.

Given this training set, the algorithm only has to try each possible partition of values that do produce an effect and values that don't produce an effect. For the social influence a, this reduces the number of values to try to a few dozen, depending mostly on the number of friends.

Applying the social and the quality threshold to an observation reduces it to one of only four possible cases (see Table 4.2). The other parameters, β and U_{min}, determine how the user will react to each of these cases. As with thresholds h and p, only a few value combinations[2] are enough to specify all possible behaviours (see Table 4.3). In fact, since the thresholds reduce each observation to two binary values, this is not a case of choosing one of the values at random because there is no more information available. The agent will behave the same with unseen data[3], whether $\beta = 0.9$ or $\beta = 1.0$.

To find an optimal solution for a given set of cases, the algorithm iterates through each of the value sets defined above. It returns those combinations of values for which more cases are classified correctly than for any other combination. In a later step, the values of β and U_{min} are analysed to classify the user according to one of the

[2] The algorithm tries the values 0.0, 0.5 and 1.0 for both β and U_{min}

[3] This is remarkable, as the model's author goes out of its way (Delre, 2007, p. 72) to randomize the value for β, which according to the reasoning presented here would have no effect at all.

x	y	Meaning
0	0	No pressure
1	0	Social influence only
0	1	Product influence only
1	1	Both social and product influence

Table 4.2: Possible values for x and y

Type	$\boldsymbol{\beta}$	$\mathbf{U_{min}}$	Description
1.	any	0.0	always adopt
2.	any	1.1	never adopt
3.	1.0	1.0	adopt when social threshold is met
4.	0.0	1.0	adopt when product threshold is met
5.	0.5	0.5	adopt when either threshold is met
6.	0.5	1.0	adopt when both thresholds are met

Table 4.3: User types, and values for β and U_{min} to produce their behaviour

types listed in Table 4.3. This characterises the influence of the social threshold and the product threshold on the user.

4.4 Calibration

As described in 2.1.2, calibration of the agent-based model can be done in two parts:

1. *Agent Specification.* Finding the agent parameters that best fit the observed behaviour.

2. *Environmental Specification.* Characterizing the network structure so that it can be accurately replicated.

These parts will be discussed in the remainder of this section.

4.4.1 Agent specification

Although agent behaviour is specified at the micro level, behaviour parameters are usually determined from macro-level data. Not surprisingly, micro-level data can provide a wealth of additional information to help choose appropriate values for these essential model parameters.

The following steps were identified in subsection 2.1.2:

Determining agent characteristics from a set of real-world cases

1. Extract cases from the data set
2. Group cases by user and decision
3. For each decision, determine the state of the environment at the time
4. Determine the parameters that best predict the decision, based on the state of the environment
5. Analyse the distribution of parameter values

Extracting cases Although identifying cases was a trivial exercise in this example model, it requires some assumptions when applying the steps to our real-world data set. The model specifies that users will hear of new songs through word of mouth or external promotion. Only after they hear of a song do they decide whether they want to listen to it. Based on this, we can define extract both 'adoption' and 'non-adoption' cases (step 1):

- *Adoption cases* describe events in which a user chose to listen to a song. These are defined as any instance in which a user listened to a song for the first time.
- *Non-adoption cases* describe events in which a user hear of a song, but chose not to listen to it. Obviously, the listening data does not specify the instances in which this happened. However, the model does specify that a user will make a decision about a song one timestep after one of his friends chose to adopt. Setting a timestep to be one day, we can define non-adoption cases as any instance in which a user had not adopted the day after a friend had listened to a song.

The model also gives agents a chance to adopt if they hear of a song through mass media. Unfortunately, there is no way to determine who was reached by mass media marketing, and we are forced to ignore those non-adoption cases. This does mean the number of non-adoption cases will be incorrect, and the *a priori* probabilities for each of the two decisions will be incorrect.

To compensate, two approaches to finding the best parameter values will be tried. The first approach uses the cases as defined above ('regular case set'), finding the parameter values that accurately predict the outcome of as many cases as possible. The second approach will first resample the cases ('resampled case set') such that there is an equal number of adoption and non-adoption cases. This will be done by duplicating a random case in the smallest category until both categories are of equal size. This has the advantage of creating equal *a priori* probabilities, removing any effect caused by the skewed distribution of adoption and non-adoption cases.

For both approaches, the steps were continued as usual and the cases were grouped by user and by action (step 2).

Defining the state of the environment The model requires two inputs to reach a decision: social pressure and the quality of the product. Unfortunately, further analysis is required to determine these values from the data set (step 3).

The model defines social pressure, a, as the proportion of friends who have previously adopted. We have the opportunity to extend this definition to include other types of relationships. In particular, it could be extend to include 'neighbours'. As chapter 3 explains, these neighbour ties are determined automatically by Last.fm algorithms, based on the similarity in taste. In this study however, the model was limited to regular friend ties, as this concept has been explored more extensively than the musical neighbours and because it links well with the structure of agent-based models. However, musical neighbours are quite an interesting element of the Last.fm concept and it would be a fascinating topic for future research.

Product quality, q, is defined in the model as an objective measure of excellence, to be compared to the user's preference p. Of course, with a product as sensitive to taste as music, product quality will be perceived differently for different users. As such, the measure of product quality must be made subjective as well, and be defined as a measure of fit to the user's taste.

The most reliable taste information made available by Last.fm is the user's listening history. Song quality is defined by considering how often[4] the user listens to that artist and how similar[5] the artist is to the the artist most listened to by the user.

For a given user and a given song, song quality is defined as:

$$q = \begin{cases} 1.0 & \text{if the artist is among the user's top 20 most listened artists,} \\ 0.8 & \text{if the artist is among the user's top 20–50 most listened artists,} \\ 0.6 & \text{if the artist is more than 80\% related to the user's top 20 artists,} \\ 0.3 & \text{if the artist is more than 50\% related to the user's top 50 artists,} \\ 0.0 & \text{otherwise.} \end{cases}$$

The numeric values of q presented here are arbitrary. However, this does not matter. Since song quality is thresholded in the model, it is merely important to ensure that higher quality corresponds to a better fit with the user's taste. Though the definition proposed here discretizes taste, it should provide a valuable approximation of the fit.

[4] As defined by the artist's place on Last.fm's individual charts for that user

[5] As defined by Last.fm's artist similarity algorithms

Type	Count	Perc.	Description
1.	314	21.4	always adopt
2.	274	18.7	never adopt
3.	36	2.5	adopt when social threshold is met
4.	665	45.3	adopt when product threshold is met
5.	97	6.6	adopt when either threshold is met
6.	81	5.5	adopt when both thresholds are met

Table 4.4: Distribution of user types (regular case set)

Type	Count	Perc.	Description
1.	0	0.0	always adopt
2.	55	3.7	never adopt
3.	19	1.3	adopt when social threshold is met
4.	1155	78.7	adopt when product threshold is met
5.	228	15.5	adopt when either threshold is met
6.	10	0.7	adopt when both thresholds are met

Table 4.5: Distribution of user types (resampled case set)

User types Using this regression algorithm, values were determined for h, p, β and U_{min} to characterise the behaviour of each of the users. The latter two parameters characterise the users as one of the six types of user listed in Table 4.3. These types indicate whether or not they appear to be influenced by their social environment and by the song. The distribution of user types among the sample population is shown in tables 4.4 (regular case set, see the explanation earlier in this chapter) and 4.5 (resampled case set).

Results for the regular case set show that for 40% of the users the social and the product influences do not improve classification. However, results for the resampled case set show that this is largely caused by the skewed a priori probabilities. Many users either have a lot more non-adoption cases than adoption cases, or vice versa.

Based on the resampled case set, the algorithm classifies nearly 80% of the users as 'adopt when product threshold is met' (type 4). This shows that product quality (fit to the user's taste) in combination with a product threshold is the best predictor for adoption for the vast majority of users.

A small group of users is classified as type 5, indicating that the user appears to adopt whenever either of the thresholds is exceeded. Further examination shows that for this group, the social and the product thresholds are almost equally valuable in correctly predicting whether the user will adopt. This suggests that social influence does play a role for at least a small group of users.

	Adoptions		Non-Adoptions		
Value	Count	Perc.	Count	Perc.	Description
1.0	43604	16.2	18639	3.0	artist in user's top 20
0.8	31698	11.8	17254	2.8	artist in user's top 50
0.6	12514	4.6	18581	3.0	artist 80%+ related to top 20
0.3	56685	21.1	103980	16.6	artist 50%+ related to top 50
0.0	124712	46.3	468857	74.7	no match with taste

Table 4.6: Distribution of product quality

Value	Count	Perc.	Description
1.0	161	11.0	artist in user's top 20
0.8	285	19.4	artist in user's top 50
0.6	170	11.6	artist more than 80% related to top 20
0.3	227	15.5	artist more than 50% related to top 50
n/a	624	42.5	not applicable for user type

Table 4.7: Distribution of product threshold, where applicable (regular case set)

Product influences Table 4.6 shows the match between the users' tastes and the songs they evaluated. This table clearly shows that songs that were adopted are more likely to have a high quality rating. This suggests a correlation between quality and probability of adoption. However, the table also indicates that the measure of product quality should be perhaps be refined. Almost half the adoptions are ranked as 'no match with taste'. This could either indicate that the users in the sample group are very adventurous and eager to try music, or simply that the measure of product quality fails to recognise how these songs fit the user's taste.

The product threshold distributions for both case sets (tables 4.7 and 4.8) show no obvious 'best' threshold. This could indicate that the adventurousness of listeners varies widely, with some only trying what they know and love and some straying more to other types of music. However, the data set is relatively small with on average only 35 adoption cases for each user, many of which come from the same album. For this reason, the eclecticity of their listening behaviour could be uncommon for them. A data set tracking listening behaviour over a longer period could provide

Value	Count	Perc.	Description
1.0	101	6.9	artist in user's top 20
0.8	372	25.4	artist in user's top 50
0.6	288	19.6	artist more than 80% related to top 20
0.3	632	43.1	artist more than 50% related to top 50
n/a	74	5.0	not applicable for user type

Table 4.8: Distribution of product threshold, where applicable (resampled case set)

Adoptions		Non-Adoptions		
Count	Perc.	Count	Perc.	Range
30083	75.8	521340	85.7	0.01 - 0.20
5374	13.5	41813	6.9	0.21 - 0.40
2362	6.0	24077	4.0	0.41 - 0.60
247	0.6	930	0.2	0.61 - 0.80
1624	4.1	20044	3.3	0.81 - 1.00

Table 4.9: Distribution of the fraction of adopters in a user's network

Count	Perc.	Range
107	7.3	0.00 - 0.20
74	5.0	0.21 - 0.40
16	1.1	0.41 - 0.60
10	0.7	0.61 - 0.80
7	0.5	0.81 - 1.00
1287		not applicable for user type

Table 4.10: Distribution of social threshold, where applicable (regular case set)

more insight into this.

Social influences Table 4.9 shows the percentage of friends that had adopted when a user chose to adopt. Inherent to the definition, this table does not include any non-adoptions where no friends had adopted. For this reason, adoptions in which no friends had previously adopted were excluded, to ensure that the numbers provided a balance overview.

Congruent with social contagion theory, the percentage of friends who had adopted previously is higher for adoptions than for non-adoptions. However, the difference is quite small, indicating that any social effect would be quite small as well.

Tables 4.10 and 4.11 show the social thresholds chosen for users with an applicable

Count	Perc.	Range
160	10.9	0.00 - 0.20
66	4.5	0.21 - 0.40
19	1.3	0.41 - 0.60
9	0.6	0.61 - 0.80
3	0.2	0.81 - 1.00
1240		not applicable for user type

Table 4.11: Distribution of social threshold, where applicable (resampled case set)

user type (type 3, 5 or 6). These results are largely unremarkable. The distribution is comparable to that of a (Table 4.9), although thresholds in the 0.21 - 0.60 range are more popular than one might expect. This could be explained by the observation that all categories of a higher than 0.21 have more adopters than non-adopters. So when more than a fifth of a user's friends has adopted, the probability of an adoption is higher than that of a non-adoption.

One remarkable observation is that almost no one has a social threshold in the 0.81 - 1.00 range. Apparently there are virtually no true 'laggards' in the sample group. Of course, this could be related to the number of distinct products on the market. With such diversity, it is a rarity that everyone in a social circle enjoys the same song.

4.4.2 Environmental specification

In an agent-based simulation model, a group of virtual agents interact to produce emergent behaviour. This behaviour depends highly on how these agents are able to communicate with each other. For example, Bonabeau (2002) experimented with a network of product-adopting agents and found that the speed of adoption was increased greatly if agents were more clustered, even if the initial user population is located entirely within one cluster. This shows that network structure can significantly affect the results, and illustrates the importance of accurately reproducing the network structure if the intention is to imitate a real-world network.

In subsection 2.1.2, the following steps were identified for analysis the structure of the network of Last.fm users:

Determining network structure in a socially connected data set

1. Calculate the degree distribution of all individuals in the data set.
2. Show this distribution in a log-log plot to visually determine the type of network, according to the characteristics listed above.
3. Use curve fitting algorithms to determine the model parameters of either a power law or an exponential distribution, depending on the network class.
4. Use existing algorithms to create the desired network class with the appropriate parameters, recreating the degree distribution.

As described in that section, the structure of the friends network can be characterised through analysing the degree distribution, the probability distribution of the number of friends (step 1). The tail of this distribution (Figure 4.1) falls on a straight line in this log-log graph, suggesting a power-law distribution (step 2).

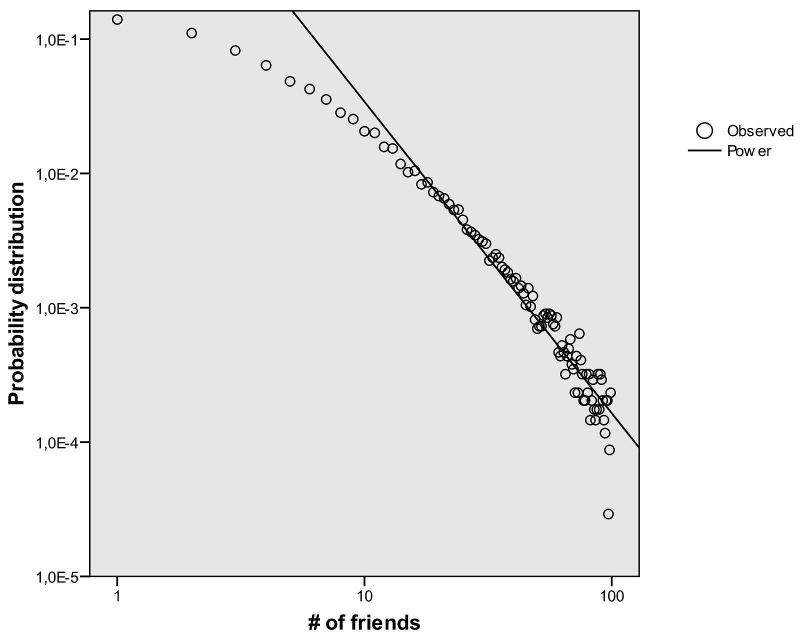

Figure 4.1: Probability distribution of the number of friends (log-log)

This kind of distribution, with a power-law cut-off, matches the characteristics of a scale-free network (Amaral et al., 2000). Many large networks have been found where the vertex connectivities follow a scale-free power-law distribution. This means that the probability $P(k)$ that a random vertex is connected to k other vertices decays for large k as $P(k) \sim k^{-\gamma}$.

As mentioned before, a scale-free network is necessarily the consequence of a constant addition of new users and a preference of new users to connect to well-connected users (Barabási & Albert, 1999). The constant addition of new users seems obvious with an open social network such as Last.fm. As for preferential attachment, it seems probable that well-connected users are more active in connecting to friends or future friends who are new to the service. With these two processes, a scale-free network would appear to be the expected structure.

The continuous line in Figure 4.1 shows the best possible power-law approximation of the vertex connectivity distribution, with $\gamma = 2.32$ (step 3). This distribution was fitted to the probability data for $k \geq 10$, with the coefficient of determination $R^2 = 0.95$ (1.0 being a perfect fit).

Together with the foundation laid down by the work of Delre (2007), the agent specification and environmental specification provide sufficient information to construct a simulation model. The next chapter will evaluate this model and discuss its suitability for examining the role of social contagion in the adoption of music.

5 Validation

This section will attempt to establish the validity of the data and model presented in the previous two chapters. In turn, it will cover:

1. data set and data source
2. model calibration
3. model validation

5.1 Data

A number of issues can be identified that could affect the usefulness of the data as a resource in researching consumer behaviour:

1. Last.fm only tracks songs played on PCs or iPods. Their software is unable to track songs listened to on CD players, radio stations, restaurants, or at a friend's house.

   For radio stations, restaurants and other places where the listener has no control over the music, this could actually be a benefit for this study. In these instances, many of the plays will not indicate an adoption, as the music was forced upon the listener. These cases are now automatically excluded from the data set.

   However, the data set also fails to include many real adoptions, because these sometimes occur in an unmonitored environment. In those cases, it seems plausible that the user will later listen to these songs on his or her own computer. In that case, the adoption itself would be recorded, but the time of adoption would be recorded incorrectly.

2. Like information from profiles, listening data can be socially biased, as users consciously choose which information they present about themselves. It can be argued that the data on the listening behaviour of Last.fm users hardly suffers from this problem. It is still possible for users to remove certain songs from their listening history, but they will have to take action to do so (opt-out) as opposed to the profiles, where the user has to choose to list a certain artist as a favourite (opt-in).

Also, there is less incentive to bias the data. Profiles serve no other purpose than presenting the user to the outside world, but an accurate listening history improves the quality of Last.fm's music recommendations. If the user consistently biases his listening history, the recommendation system will soon become useless to him, creating a strong incentive not to do so. This makes it much more likely that the data is a direct representation of the users' behaviour.

3. The friends identified on the Last.fm site only include those friends who use the Last.fm service. No doubt, in almost every case a user's social network will be larger than this group. This means that the data set will not show many of the social interactions that caused an adoption. In reality, any social effect on adoption of music can be much larger than is reflected in this data set.

4. It is assumed that only adoptions spread through the social network. However, it may be possible that non-adoptions spread as well. 'Hey, look at this really bad song I found!'

5. The users included in the sample group were chosen by selecting a sample group from the members of a Last.fm 'group' and crawling these to find connections to the first and second degree. Because there is a not a complete list of Last.fm users available, it was not possible to get a more representative group of users. Still, it could be questioned whether the listening behaviour found in the data set is representative for the listening behaviour of all Last.fm users, let alone that of non-users.

   However, the listening data in the data set was compared to the listening data for all Last.fm users found in the top 400 charts. At least for hit songs, this comparison shows that the growth in listener numbers in the data collected here is comparable to the growth in global listener numbers.

   The relation between global listener numbers and local listener numbers differs per track – for the average Coldplay track the data set contains the behaviour of 10% of all Last.fm users who listen to that track, as opposed to only 6% for the average Madonna song – but is generally quite consistent between tracks of the same artist as well as through time[1] (see figures 5.1, 5.2 and 5.3). Of course, this still leaves open the question of this behaviour being representative for non-users as well.

This section has listed some possible issues with using Last.fm as a data source. Some of these warrant further inspection (see chapter 8 for suggestions on future research) but none seem serious enough to completely discard the data set. It will suffice at least for current purposes: demonstrating the usefulness of micro-level data in agent-based modelling.

[1] The percentage of all users included in the data set varies usually by less than a percent in consecutive weeks.

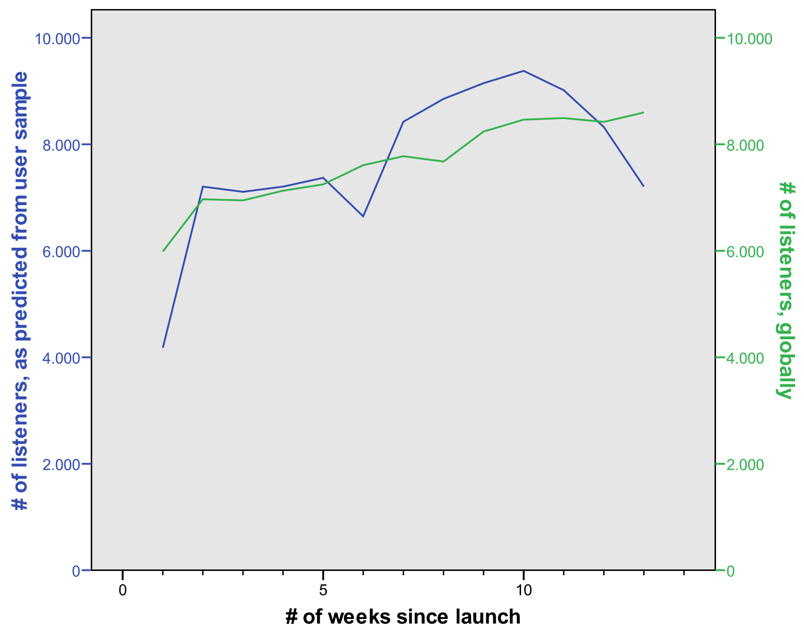

Figure 5.1: Global and local listener numbers for Paramore's 'That's What You Get'

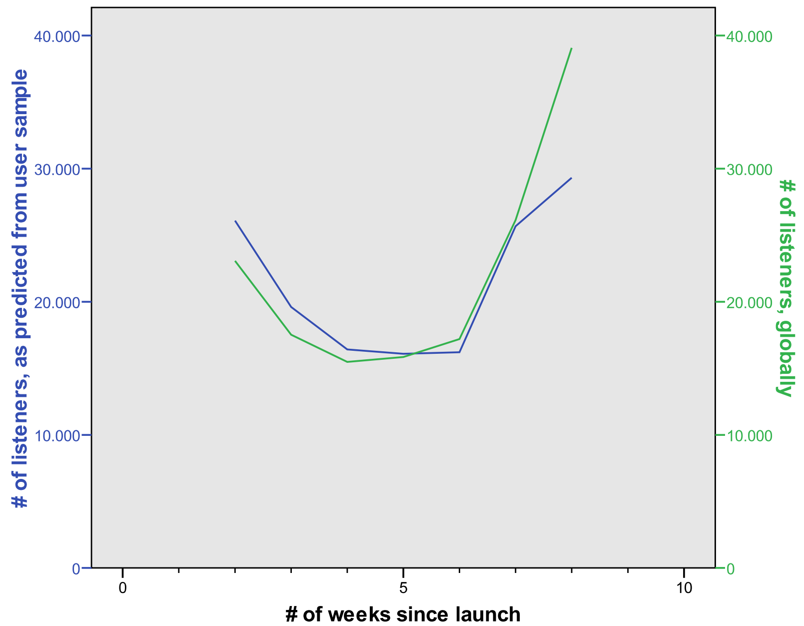

Figure 5.2: Global and local listener numbers for Coldplay's 'Violet Hill'

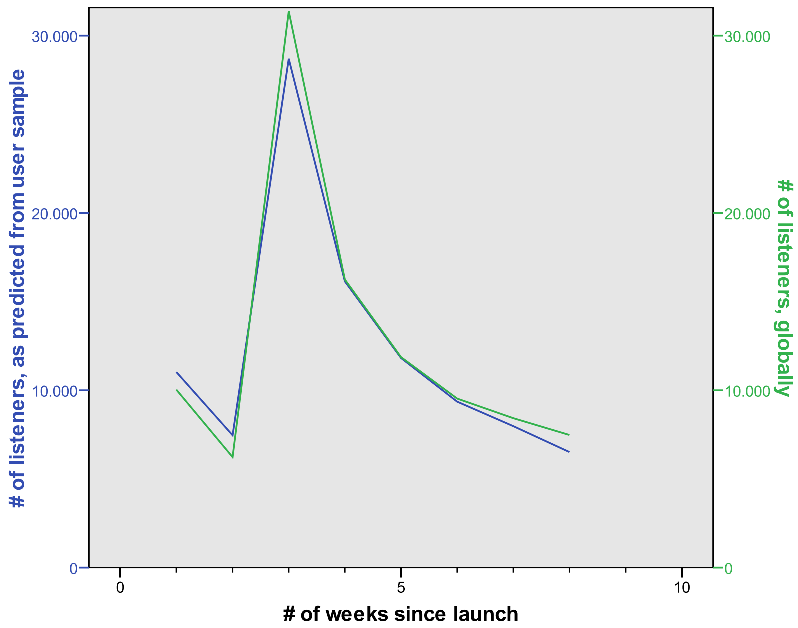

Figure 5.3: Global and local listener numbers for Nine Inch Nails' 'Discipline'

5.2 Calibration

Validating the model at a micro level is closely related to the regression steps described in the previous section. The model's precision in predicting an agent's decision is a good indicator of the predictive value of the environment inputs. As such, it can signal whether essential processes or signals are missing from the model.

Tables 5.1 and 5.2 show for respectively the regular and the resampled case sets the precision for each user type. One noticable difference is that the precision on adoption cases is higher in the resampled case set. This is to be expected, as adoption cases are generally in the minority, so the resampled case set will attach more weight to maximizing precision on adoption cases. Note that the value of calculating precision on a separate test set becomes apparent here. It would be interesting to see the difference in performance between the regular and the resampled case sets on a different test set. Failing this, the interested reader is referred to the section on simulation validation, where the performance of both case sets is tested with simulations on a separate evaluation set.

Looking at the precision figures for each of the user types, it is interesting to note that types 3 (social-only) and 6 (both social and product) have a relatively low precision on adoptions. Although social influence is an important factor in predicting adoption for this group, this suggests that it is more valuable in predicting non-adoptions than it is in predicting adoptions. Apparently, these users are more likely to not adopt if

User type	Precision	
	Total	Adopt
1. always-adopt	0.75	1.00
2. never-adopt	0.91	0.00
3. social-only	0.91	0.17
4. product-only	0.80	0.45
5. either	0.76	0.48
6. both	0.91	0.14
Total	0.80	0.55

Table 5.1: Model precision, for each of the user types (regular case set)

User type	Precision	
	Total	Adopt
1. always-adopt	n/a	n/a
2. never-adopt	0.50	0.00
3. social-only	0.61	0.46
4. product-only	0.70	0.55
5. either	0.70	0.57
6. both	0.82	0.67
Total	0.66	0.65

Table 5.2: Model precision, for each of the user types (resampled case set)

their friends had not adopted than they are to adopt if their friends had listened to the song first. This could suggest that many of their friends have similar taste and they rarely listen to music that their friends hadn't listened to. However, this would also mean that this group consist of laggards, who are slow to adopt new music. Either way, the low precision on adoption cases questions the importance of friends in finding new music.

Total precision is highest in the regular case set, with 80% of the cases classified correctly. This is significantly larger than classification by a random coin toss. Even the resampled case set, which does not have the benefit of the a priori probability, outperforms random classification with 66% precision. This suggests that the social and product influences do have predictive value in determining whether a consumer will adopt. However, there does seem to be space for improvement.

5.3 Simulation

Section 2.1.3 described the history-friendly approach to validation: choosing the parameter values that best approximate observed history. It identified the following steps:

Validation, using the history-friendly approach

1. Implement the agent-based model
2. Initialise agents and environment
3. Simulate adoption of all products in the test set
4. Repeat steps 2 & 3 several times to filter random deviations
5. Compare simulated adoption figures to real data

The implementation of the simulation software (step 1) was preceded by a survey of available agent-based modelling tools. This survey revealed that many of the popular tools are not equipped to process large sets of micro-level data for comparison. Because of this, the model was implemented using custom software, built to integrate with the existing database server, but architected in a modular fashion to allow for different network and agent initialization algorithms to be compared.

As subsection 2.1.3 explained, not all possible parameter values can be evaluated through simulation because of the computational complexity. Instead, calibration was done based on the fit to the training data (see section 4.4) and validation would be performed on a limited number of candidates.

Candidates were explored in three dimensions:

1. Network structure:
   a) random: each agent connects to one randomly chosen other agent to which it was not already connected,
   b) circle: a circle network, with each agent connected to exactly two other agents, forming a single circle,
   c) scale-free: a network is generated with a scale-free structure, matching the structure found in section 4.4, or
   d) identical: a direct copy is made of the network structure described by the data set
2. Parameter set:
   a) regular: a random sample from a Normal distribution with the mean and standard deviation of the parameter values determined for the regular case set,
   b) resampled: the same, for the resampled case set, or

c) identical: for each agent, all parameters are set equal to the parameters found for their real-life counterpart[2] or, if none are available, using the method for 'regular'.

3. U_{min} and β:

a) normal: initialize U_{min} and β like any other parameter, using a Normal distribution based on the parameters found in the case set, or

b) typedist: initialize through a type distribution, assigning user types according to Table 4.4 or 4.5 and setting U_{min} and β according to Table 4.3.

Each possible combination of these $((4 \cdot 2 \cdot 2) + 2 =)$ 18 candidates was explored (step 2). For each candidate, simulations were run (step 3) on all products in the test set (3743 products). The number of iterations was set equal to the number of days for which comparative data was available. Each simulation was repeated 10 to 50 times (step 4), depending on how long it took for the diffusion curve to stabilise. Initialized networks of agents were reused for each product, but never more than once for one product.

An obvious ways to evaluate the similarity between simulated and real-world adoption is to calculate the coefficient of determination, R^2. However, even if this produces an unsatisfactory low value, there are other ways to explore the relationship. If a linear of even non-linear mapping could be found that show a satisfactory relationshop between the simulated and the real data, this would reveal a bias in the model affecting its performance. For example, if the model always predicted adoption to be twice as high as the real adoption, a linear mapping ($f(x) = 2 \cdot x$) would show this relation.

Although this method of analysing the fit might appear to be stretching the results to find a relationship, there are two major benefits. First of all, revealing a bias in the model could provide information that is very valuable in refining the model. Secondly, even if an additional calculation step (applying the linear or non-linear mapping) was needed to correct the results, this would not diminish the model's usefulness. Despite the bias, the model could still be valuable in predicting or understanding the dynamics of the market that is being described.

Simulation results are listed in tables 5.3 and 5.4. A few remarkable things can be observed here:

1. The performance of the scale-free network structure is very close in performance to the 'identical' network structure, which was a one-to-one copy of the original network. This suggests that a scale-free network is indeed a good approximation of the network structure of these Last.fm users.

[2]This means that the identical parameter set can only be used with the identical network structure, otherwise there is no real-life counterpart

Network structure	Parameter set	U_{min} and β	# of simulations	Corr.	R^2 (linear)	R^2 (non-linear)
circle	regular	normal	41,130 runs	0.56	0.31	0.36
circle	regular	typedist	44,318 runs	0.59	0.34	0.39
circle	resampled	normal	41,031 runs	0.55	0.30	0.38
circle	resampled	typedist	40,116 runs	0.46	0.21	0.30
identical	identical	normal	101,454 runs	0.30	0.09	0.23
identical	identical	typedist	89,575 runs	0.38	0.14	0.27
identical	regular	normal	67,006 runs	0.28	0.08	0.22
identical	regular	typedist	186,830 runs	0.36	0.13	0.52
identical	resampled	normal	88,573 runs	0.29	0.08	0.23
identical	resampled	typedist	87,716 runs	0.47	0.22	0.32
random	regular	normal	41,212 runs	0.57	0.33	0.37
random	regular	typedist	46,074 runs	0.59	0.35	0.37
random	resampled	normal	41,552 runs	0.62	0.39	0.43
random	resampled	typedist	40,780 runs	0.50	0.25	0.31
scale-free	regular	normal	174,812 runs	0.36	0.13	0.36
scale-free	regular	typedist	186,830 runs	0.36	0.13	0.55
scale-free	resampled	normal	111,407 runs	0.32	0.10	0.25
scale-free	resampled	typedist	128,580 runs	0.47	0.22	0.33

Table 5.3: Results of 1,558,996 simulations (candidates sorted by parameters)

	Network structure	Parameter set	U_{min} and β	# of simulations	Corr.	R^2 (linear)	R^2 (non-linear)
1.	scale-free	regular	typedist	186,830 runs	0.36	0.13	0.55
2.	identical	regular	typedist	186,830 runs	0.36	0.13	0.52
3.	random	resampled	typedist	41,552 runs	0.62	0.39	0.43
4.	circle	regular	typedist	44,318 runs	0.59	0.34	0.39
5.	circle	resampled	normal	41,031 runs	0.55	0.30	0.38
6.	random	regular	typedist	46,074 runs	0.59	0.35	0.37
7.	random	regular	normal	41,212 runs	0.57	0.33	0.37
8.	scale-free	regular	normal	174,812 runs	0.36	0.13	0.36
9.	circle	regular	normal	41,130 runs	0.56	0.31	0.36
10.	scale-free	resampled	typedist	128,580 runs	0.47	0.22	0.33
11.	identical	resampled	typedist	87,716 runs	0.47	0.22	0.32
12.	random	resampled	typedist	40,780 runs	0.50	0.25	0.31
13.	circle	resampled	typedist	40,116 runs	0.46	0.21	0.30
14.	identical	identical	typedist	89,575 runs	0.38	0.14	0.27
15.	scale-free	resampled	normal	111,407 runs	0.32	0.10	0.25
16.	identical	identical	normal	101,454 runs	0.30	0.09	0.23
17.	identical	resampled	normal	88,573 runs	0.29	0.08	0.23
18.	identical	regular	normal	67,006 runs	0.28	0.08	0.22

Table 5.4: Results of 1,558,996 simulations (candidates sorted by maximum R^2)

2. Even though the performance is very similar, 'scale-free' performs slightly better than the 'identical' network structure in almost all cases. This applies to all parameter sets and for both correlation, R^2 with a linear model and R^2 with a non-linear model.

3. The candidates with network structure and agent parameters directly copied from the original network ('identical' and 'identical') rank 14^{th} and 16^{th} in comparison to the other candidates, almost at the bottom of the list. Note that this is the closest approximation of the original network structure and distribution of agent parameters. Apparantly, generalizations of the network structure and parameter distribution are better at predicting adoption than a model that is exactly fitted to the original data.

4. Candidates based on the regular case set generally outperform candidates based on the resampled case set. This suggests that the a priori probabilities for adoption/non-adoption are better approximated in the regular case set than in the resampled case set. In the resampled case set, cases were resampled at a per user basis to create equal a priori probabilities (0.5/0.5).

5. Candidates using a type distribion for U_{min} and β generally outperform candidates using a Normal distribution to assign those parameters. This suggests that this type distribution is a better way to represent the heterogeneity in user types.

6. Candidates with a circle or random network structure perform surprisingly well. These simple network structures were expected to perform very badly, because the networks generate are so much different from the original networks. Their surprisingly good performance can probably be attributed to the fact that they generally underpredict adoption. Since most songs are adopted by very few people, the difference between a very low number and the real data will generally be smaller than the difference between a very large number and the real data.

The best candidate uses a scale-free network structure, the regular parameter set and a type distribution to assign U_{min} and β. For this candidate, the R^2 with a non-linear model (in this case, an exponential model) is much higher than the R^2 with a linear model, suggesting that this model dramatically overpredicts adoption for songs that do well. It achieves a R^2 of 0.55 with a non-linear model, indicating that 55% of the variability in the data is explained by this model. Certainly room for improvement, but more than good enough for an exploration of this domain in the current state of research on this topic.

6 Results

The previous chapters described three parts of constructing an agent-based model. Chapter 3 described how the music choices of a group of Last.fm users were recorded. Chapter 4 described how this data was used to build an agent-based model of the adoption of songs. Finally, chapter 5 reflected on the data set and model and examined potential issues with either of these. No objections were found that made either the data set or the model unsuitable for examining the role of social contagion in the adoption of music.

This chapter will apply the model and data set to answer the research questions identified in chapter 1. In turn, it will examine the role of social contagion in the adoption of songs, explore the characteristics of influentials and imitators, and discuss the consequences of having many consumers who are influenced by their social circle.

6.1 Social contagion in song adoption

This thesis set out to examine whether Last.fm users are affected by their friends when deciding what music to listen to, i.e. whether social contagion plays a role in the adoption of music. Earlier in this report, section 4.4 analysed which agent parameters most accurately reproduce the music choices observed and recorded in the data set. These parameters provide insight into the factors that correlate with the decision to adopt.

For the majority of users, only the match with taste was valuable in predicting adoption. This suggests that for this group social contagion plays a very small role in the decision to adopt. It is quite possible that the role of social contagion is too small to detect in this manner. Fortunately, there is another measure that could provide insight into its importance: the q/p ratio.

The q/p ratio refers to the q and p in Bass' equation for the rate of adoption (see chapter 4). In this equation, q is a measure of social influence and p is a measure of product influences. Thus, the q/p ratio provides a measure of the importance of social influence in a given market.

The general approach to calculating the q/p ratio is to fit the rate of adoption to Bass' equation ($r(t) = p + qF(t)$, see Bass, 1969). However, this is only applicable

	q	p	q/p
adoption = listening 1 time	15232	90062	0.17
··· songs with 51-100 listeners	2535	37155	0.07
··· songs with 101-250 listeners	3740	30814	0.12
··· songs with 251-500 listeners	2314	9030	0.26
··· songs with 501+ listeners	6643	13063	0.51
adoption = listening 2 times	7354	33746	0.22
··· songs with 51-100 listeners	838	10845	0.08
··· songs with 101-250 listeners	1443	11350	0.13
··· songs with 251-500 listeners	1099	3993	0.28
··· songs with 501+ listeners	3974	7558	0.53
adoption = listening 6 times	1927	7056	0.27
··· songs with 51-100 listeners	164	1567	0.10
··· songs with 101-250 listeners	290	1567	0.10
··· songs with 251-500 listeners	241	935	0.26
··· songs with 501+ listeners	1232	2472	0.50
adoption = listening 10 times	948	3206	0.30
adoption = listening 15 times	524	1641	0.32
adoption = listening 20 times	336	1023	0.33

Table 6.1: Overview of q/p ratios for different definitions of 'adoption' and varying degrees of popularity.

when analysing a single diffusion curve. There is no commonly used definition for an aggregate q/p ratio, for a collection of diffusion curves. Jager (2008)'s definition is used here: q/p equals the number of adoptions with social influence divided by the number of adoptions without social influence. Here, an adoption with social influence, is an adoption that was preceded by an adoption of a friend or connection.

Table 6.1 lists the q/p ratio, for different definitions of 'adoption' and distinguishing between the q/p ratio for niche songs (few listeners) and the q/p ratio for mainstream music (many listeners). The results shows that the q/p ratio is very small, regardless of the definition. The ratio hardly exceeds 0.50, while a survey of comparable studies showed an average ratio of $q/p = 30.57$ (van den Bulte & Stremersch, 2004)[1]. However, products (songs) in the music industry are not durables and a low q/p ratio – signalling very low social influence – is to be expected (van den Bulte, 2008) considering the incredible number of products available.

[1] van den Bulte & Stremersch (2004) surveyed 52 consumer durables in 46 publications, and found q/p ratios with $\mu = 30.57$, $\sigma = 8.41$, min = 0.00036, max = 57526.44

	Influentials	Imitators
0 - 1 friends	6.5%	24.0%
2 - 5 friends	32.2%	38.0%
6 - 10 friends	21.6%	13.9%
11 - 15 friends	13.6%	7.3%
16+ friends	26.0%	16.7%

Table 6.2: Number of friends for influentials and imitators

6.2 Influentials and Imitators

A popular theory in the field of diffusion research states that the success of a new product can be heavily influenced by a small number of 'influentials'. If these consumers adopt, their friends will be more inclined to follow them. This second group, users who are influenced by their social circle, is usually referred to as 'imitators'.

Influentials Influentials, the theory claims, generally have a large number of friends (large degree centrality) or are in some other way well-connected (large eigenvector centrality, or some other measure). To test this assumption using the micro-level data set, a measure of 'influence' was calculated for each adoption. Influence was defined as the number of friends who adopted on a later date (within recorded history). No correlation was found between the amount of influence and several possible user characteristics. There was a correlation between influence and the number of friends, as the assumption claims, but further study showed that the cause for this was that a larger group of friends will on average contain more adopters.

The study was hindered by the fact that virtually all products had a very low number of adopters. Because of the diversity of supply, for most users there will be music that fits their niche tastes (the famous Long Tail effect). Consequently, if any influentials existed, their influence would most likely cause two of their friends to adopt when you might predict only one to adopt. Obviously, such an effect would be too small to discard chance as a possible cause.

There seems to be no correlation between the strength of the influence and any of the user characteristics, including the number of friends. However, a different method of analysis did reveal the influence of well-connected people. For each user, it was calculated what percentage of adoption cases had an influence that was higher than would be expected considering the number of friends[2]. The results showed that users for whom this was the case for more than 9% of their adoptions, more commonly had a large number of friends (see Table 6.2).

[2] That is, higher than the number of friends who have not yet adopted times the probability that a random user adopts.

	Number of cases							
	≤ 100		101-200		201-400		401+	
Type	Count	Perc.	Count	Perc.	Count	Perc.	Count	Perc.
1.	0	0.0	0	0.0	0	0.0	0	0.0
2.	16	9.6	19	5.3	15	3.2	5	1.1
3.	1	0.6	3	0.8	8	1.7	7	1.5
4.	147	88.0	311	87.1	369	78.8	328	69.1
5.	2	1.2	22	6.2	73	15.6	131	27.6
6.	1	0.6	2	0.6	3	0.6	4	0.8

Table 6.3: Distribution of user types, split up by the number of cases for that user (resampled dataset)

Imitators Analysing the model parameters (subsection 4.4.1) gave several indications of a social effect in the adoption of new music. In particular, Table 4.5 showed that nearly 11% of the users were influenced both by the product and by the social environment (type 5).

Table 6.3 shows that users of this type are generally users for which a large number of cases is available. This means that a lot of adoptions and non-adoptions have been recorded, suggesting the user has a large number of friends or is otherwise very active on Last.fm. If valid, this observation could lead to interesting conclusions about the effects of such social networking sites on the adoption of new music. It would appear that the social effect is much stronger with heavy users of such services.

6.3 The effect of social contagion

As previous results have shown, it is possible that the influence of social contagion will increase with the use of social networking sites. The question remains: how will this affect the adoption of new products? This is the type of problem that is commonly explored using simulation modeling.

Section 2.2.3 presented the steps to simulate the effect of social contagion:

Investigating the effect of social contagion through simulations

1. Implement the agent-based model
2. Initialise agents and environment
3. Simulate adoption of all products in the test set
4. Repeat steps 2 & 3 several times to filter random deviations
5. Repeat steps 2-4 with different social contagion settings and compare results

	Variant	After 1 year
1.	100% either social or product	1680.77 adopters
2.	100% social	1322.61 adopters
3.	25% either social or product, 25% product only	1067.36 adopters
4.	75% social, 25% product	831.03 adopters
5.	50% either social or product, 50% product only	565.97 adopters
6.	50% social, 50% product	427.58 adopters
7.	25% either social or product, 75% product only	253.24 adopters
8.	25% social, 75% product	213.56 adopters
9.	100% product	103.54 adopters

Table 6.4: Simulation results (sorted by average number of adopters)

Table 6.4 lists the different mixes of user types that were simulated. These range from 100% product-focused users to 100% socially focused users or 100% users triggered by either the product or the social threshold. The results are presented in the same table, as well as graphically in Figure 6.1.

Several remarkable results can be observed here:

1. All adoption curves stabilise quickly. By day 100, the number of adopters have reached their maximum for all simulation variants. The reason is that external promotion (e_2) is only simulated for as many days as there is data available. At the time of simulation, the maximum number of days for which data was available was 94. From the last day of data to the end of the year, the song spreads purely on word of mouth, without any additional seeds. As the simulation shows, word of mouth alone is not sufficient to keep adoption going.

2. Not surprisingly, simulations where part of the population adopts when either the product or the social threshold are exceeded outperform simulations where only the social threshold is considered. Still, 'either' simulations have only 1.3x as many adopters. One might expect a bigger difference, considering that adoption in 'either' simulations will go much quicker in parts of the network that have low penetration.

3. When replacing 25% of product-focused users with 25% socially focused users, the number of adopters is multiplied by two. This multiplication is not quite so consistent for 'either' simulations: when replacing 25% of product-focused users with 25% users that adopt when either the product or the social threshold is met, the number adopters is multiplied by 1.5x-2.5x. The multiplication factor decreases: the difference is 2.5x for 0% to 25% and 1.5x for 75% to 100%.

4. There is a major difference between adoption with 100% product-focused users and adoption with 100% social-focused users. On average, adoption with 100% social-focused users leads to 13 times as many adopters. The high number of

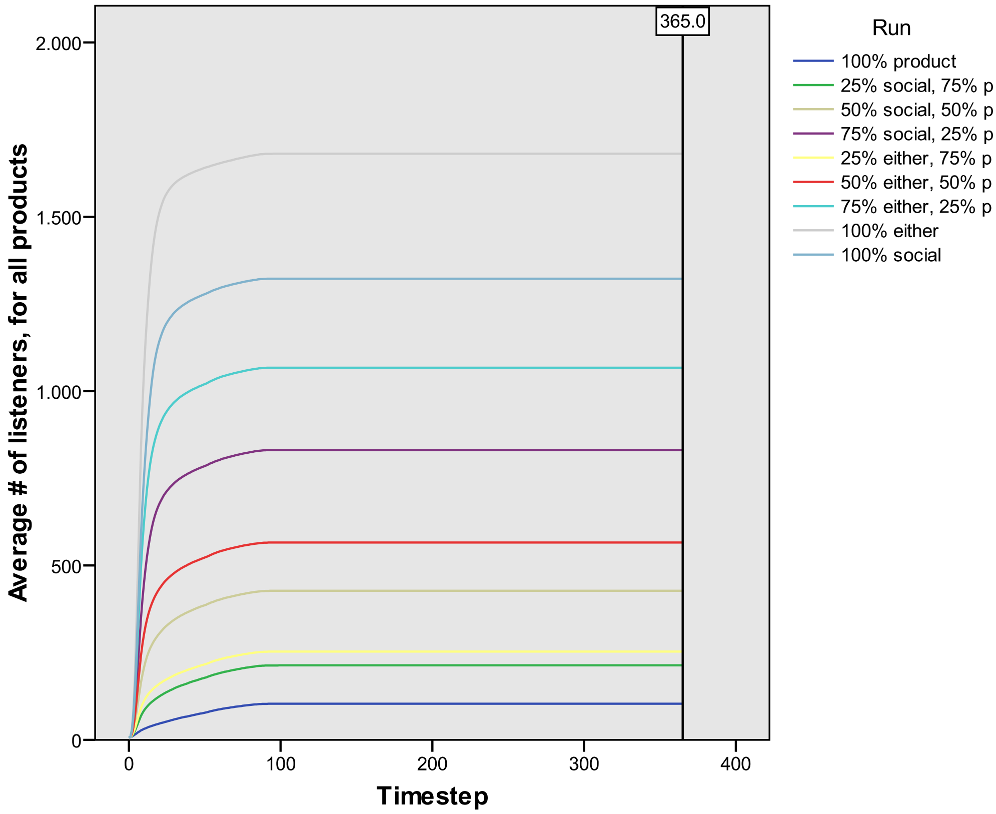

Figure 6.1: Mean adoption for varying degrees of social and product influences

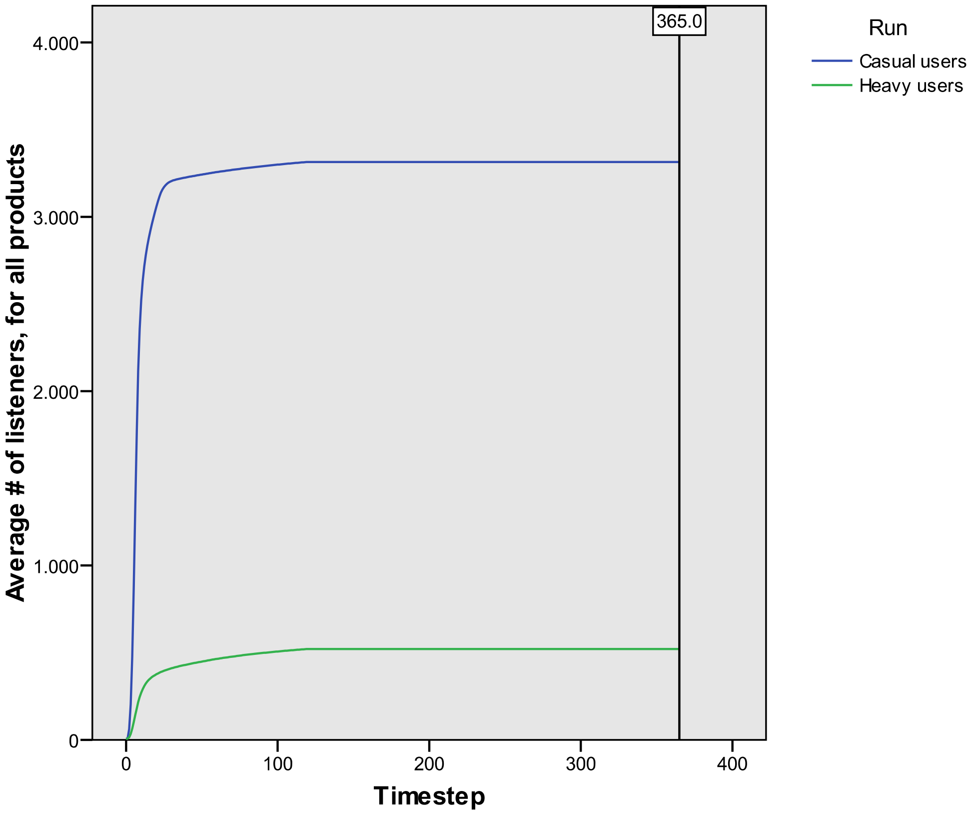

Figure 6.2: Mean adoption for varying degrees of social and product influences

adopters for the social-focused simulation is surprising, as these simulations are completely dependent on the adopters in timestep 0 to start adoption. Apart from these initial adopters (who adopt regardless of whether the social or product thresholds are met), all users wait for at least one friend to adopt first. Of course, the social thresholds used here (see 4.4) are quite low. On average, only 19% of friends need to adopt first, so often only one adopter suffices. Apparantly, with social thresholds as low as this, adoption figures can reach these levels, even if no other user adopts based on the product itself.

The results show that social contagion greatly increases the number of adopters. This suggests that if social network sites continue to grow, new songs will be able to reach larger audiences than before. Most likely this will prove especially beneficiary for independent artists, who always had trouble reaching large audiences through traditional channels. This may in the future reduce the necessity of signing contracts with major music labels, and clear the way for new musical talent and diversity.

Some caution is appropriate here, however. Additional simulation runs were done (see Figure 6.2), comparing the group of users who frequently use Last.fm to the group of users who use Last.fm very infrequently. Contrary to the earlier simulations, which tested several 'clean' mixes of social and non-users users, in these simulations the runs with social high-intensity users consistently had lower adoption than the product-focused low-intensity users. Perhaps this difference is caused by the relatively high percentage of users who always adopt or never adopt in these groups.

7 Conclusion

In the last few decades, advances in technology have caused a surge of interest in agent-based models (see chapter 1). Increasingly, the methodology has been used to work on the many challenges of the social science. ABMs are usually trained and validated on macro-level data, which describes the behaviour at an aggregate level. However, this approach often requires many assumptions and it is impossible to show or even to make a credible case that a model is an accurate representation of reality. Even though a model might generate the appropriate behaviour at a macro-level, there could be several micro-level specifications that do so, not all of which are correct.

Increasingly, there is a new type of data available that could benefit researchers who are training, validating and applying agent-based models. This type of data, micro-level data, describes behaviour at an individual level, as opposed to the aggregate statistics commonly used up to this point. As such, it provides information about how individuals react to various situations. By describing these situations, a researcher can trace back which information was available at the time of the action, and determine which features affect the decision taken.

This thesis explored one such micro-level data set, on the music choice of a group of Last.fm users. This data set describes in detail which songs each user chose to listen to, at what time they first listened to the song, which of their friends had already listened to the song and how well the song matches with what they usually listen to. This data was used to examine the processes underlying the adoption of new songs by music consumers (see chapter 2).

Building an agent-based model

Over a period of four months, 34.325 Last.fm users were monitored (see chapter 3). During this period, all the songs that they listened to were stored for analysis. Additionally, for each user their complete list of friends and musical neighbours (people with very similar taste) was stored, as well as some information about their musical taste.

During these four months (March 18 - July 18), the users in the sample group listened to close to six million tracks, well over a hundred million times in total. Out of these six million, about eighteen thousand tracks were newly released during these four months and were used for analysis. All other tracks were discarded, because it was impossible to determine whether the first time that a user listened to that song was really the first. For these eighteen thousand tracks, close to ten million plays were recorded. These plays formed the basis for the calibration of an agent-based model for the adoption of new music.

Delre (2007)'s model for the diffusion of innovation was used to model the recorded listening choices (see chapter 4). This model was calibrated on the Last.fm data set, in two parts:

1. The *agent specification*, with the parameters governing the agents' behaviour. For each user, a regression algorithm determines the parameters that best match the music choices recorded in the training set (a random two-thirds of newly released songs in the data set). An analysis of how much influence adoption by friends and match with music taste have gave insight into the types of users.

   This showed that for the majority of users, the match with music taste was the best predictor of the decision whether to adopt. For about 10% of the users both adoption by friends and the match with music taste were good predictors of adoption. Thresholds for the adoption by friends were generally low, because most songs only reach a small number of listeners, reflecting the 'long tail' of the music industry.

2. The *environmental specification*, with an analysis of the network structure. This analysis examined how many friends each user in the user sample has. The results suggest a scale-free network, with $\gamma = 2.32$.

Together with the foundation laid down by the work of Delre, the agent specification and environmental specification provide sufficient information to construct a simulation model of the adoption of music.

Before using this model to answer the research questions identified in chapter 1, both the data, the calibration results and the model itself were validated (see chapter 5):

1. For all songs that appear in the Last.fm Weekly Top 400, weekly listener counts in the user sample were compared to overall listener counts listed on the Last.fm charts. This showed that, at least for these hit songs, the recorded adoption curves are representative for the shape of the adoption curves for all Last.fm listeners. This indicates that adoption behaviour recorded in the data set is a valid sample of the behaviour of all Last.fm users.

   Several reasons were explored why data collected from Last.fm might not be representative of the music listening behaviour of the average music consumer. Some of these warrant further research, but at the moment it seems plausible that this data could at least provide insight into the choices of young internet-savvy music consumers.

2. The quality of the model calibration was assessed by determining how well the model was able to classify adoptions and non-adoptions in the training set. To do so, the model was asked to predict whether a user would adopt a song, given only the user parameters, the number of friends who had listened to the song before and the match with the user's taste. As such, the only information that was available was the information that would be available in a simulation. Based on this information, the model with the best parameter set managed to classify 80% of the cases correctly: in 80% of recorded situations, the model was able to correctly predict whether a user would listen to a song or not.

3. The validity of the model in simulation was assessed with simulation runs. Eighteen different candidate models were evaluated, to be able to examine assumptions about network structure and the distribution of agent parameters. For each candidate, 10 to 50 simulation runs were done for all songs in the test set (the remaining one-third of newly released songs in the data set) and the mean adoption figures per day were compared to actual adoption data. The best candidate approximated the test set with $R^2 = 0.55$, albeit using a non-linear mapping function to compensate for consistent overprediction.

Although the validation steps produced some directions for future researcher, they did not succeed in proving the model invalid. As such, the model can now be used to examine the role of social contagion in the music choice of Last.fm users.

Examining social contagion

Using the Last.fm data set and the model that was constructed from it, three research questions were explored:

1. What is the role of social contagion in the adoption of new songs by users of the Last.fm network?

2. Who are the influencers and the imitators on the Last.fm network?

3. How does social contagion change the adoption of new songs by users of the Last.fm network?

For most music consumers, social contagion plays a very small role in their music choice. Looking at the distribution of user types in the trained model, 90% of the agents only considers the product threshold when making their decision (see section 6.1. Only 10% of the users considers both the product and the social thresholds. This finding is confirmed by the q/p ratio, which shows that social influence in this music market is very low. Both results are to be expected considering the incredible number of songs available on the market (van den Bulte, 2008).

An interesting follow-up question is whether that 10% group is homogeneous, or if some users are easier to influence or influence more friends. This ties in to the influentials theory, which divides consumers into *influentials* and *imitators* (see section 6.2).

The search for influentials proved difficult, because of the limited social influence in this market. Still, it turned out that influentials are more likely to have many friends. This analysis already takes into account the fact that a larger group of friends (or any group of people) will on average contain more adopters.

Imitators were easier to identify. The user type distribution of users for which few cases were available was compared to that of users for which many cases were available. This comparison showed that the first group had almost no users who were influenced by their friends, while this percentage was significantly higher in the second group. Since it seems likely that the number of cases is related to the intensity of Last.fm use, this suggests that heavy Last.fm users are more likely to be influenced by their friends, and could be considered as 'imitators'.

Increased social contagion could dramatically increase adoption of new songs. Simulations were done with different mixes of product-focused consumers and socially focused consumers (see section 6.3). These showed that a group of socially focused consumers greatly accelerated adoption. Changing 25% of agents from product-focused to socially focused agents doubled adoption every time, be it from 25% to 50% social or from 50% to 75% social. The same effect (though with more variance, 1.5x-2.5x) was found when making the same number of product-focused agents also responsive to social contagion.

If the rise of social networks indeed makes more people more responsive to the music choices of their friends, these simulation results suggest that it will become easier for bands to reach a large audience. Considering that artists signed by major labels already managed to reach a large number of people thanks to mass-media, this could have an effect on independent artists in particular.

Some caution is needed here. Additional simulation runs were done, comparing the group of users who frequently use Last.fm to the group of users who use Last.fm very infrequently. Contrary to the earlier simulations, which tested several 'clean' mixes of social and non-users users, in these simulations the runs with social high-intensity users consistently had lower adoption than the product-focused low-intensity users. Perhaps this difference is caused by the relatively high percentage of users who always adopt or never adopt in these groups.

In any case, this suggests a fascinating direction for future research. The results presented here suggest that the social interaction facilitated by social networking sites does have an effect on the adoption of music, and perhaps other products. This could make consumers more susceptible to the consumer choices of their social circle, but also increase the effect of word-of-mouth marketing. In combination with other types of research, different agent-based models could help understand these changing dynamics and give insight into the changes introduced by these sites. Perhaps these changes even go so far as to change the world economy...

8 Discussion

This chapter will discuss possible limitations to the study, directions for future research and opportunities for applying agent-based modelling in combination with detailed data sets such as the one presented in this thesis.

8.1 Limitations

There are several issues with the conclusions presented in the previous chapter, which could have affected the validity of the results:

1. *Not all assumptions have been tested.* Although the study tried to explore different alternatives as much as possible, this was not always possible due to time constraints. In particular, it could have been valuable to explore alternative definitions for:

   - a timestep: a non-adoption was defined as any instance in which the user did not adopt a song *one day* after one of his friends had. Would the results have been different if it had been one hour, six hours or a week?
   - an adoption: an adoption was defined as any instance in which a user listened to a song once for the first time. Would the results have been different if adoptions were only counted after listening to a song at least two or five times? What about if the definition of an adoption was different when calculating the percentage of friends who had already adopted?
   - a friend: two users were considered friends if they had marked each other as 'friends' on Last.fm. However, Last.fm also introduced the concept of 'musical neighbours'. These neighbours are identified automatically by Last.fm's algorithms, and signal a close similarity in taste. Would the results have been different if musical neighbours were also considered friends?

2. *The data set records the adoption of new songs over a period of a few months.* It could be argued that music has a short lifespan, but it should still be noted that many diffusion studies cover a longer period. To illustrate, Bass (1969) studied the adoption of consumer durables, over a period of 14 years. Would the results have been different if data was collected over such a long period?

3. *The validity of the data set depends heavily on the validity of using Last.fm data for this type of research.* Section 5.1 lists some concerns about the data, the most important of which is: are Last.fm users representative for the average music consumer? Would the results have been different if a different group of consumers had been observed?

4. *Music from other genres could have different types of market.* Just like Last.fm users might be unrepresentative for other music consumers, the same could apply to the music genre. Although the songs tracked in the data set were not limited to any specific genre, there is most likely a bias towards Britpop or independent pop music due to the way the users were sampled. Would the results have been different if more music from other genres was examined?

The following section will outline directions for future research that could potentially address the issues identified here.

8.2 Future research

This section will reiterate through the issues identified in the previous section, and explore possible directions for future research to explore both the consequences and possible solutions to these issues. As with the first section, it will cover the issues with the agent-based model for music consumers separately from the methodology itself.

There are several ways in which the model for consumer behaviour in the music industry could be explored further:

1. *Explore the assumptions identified in the previous section.* This includes experimenting with different timestep lengths, different definitions for 'adoption' and including musical neighbours in the friends circle.

2. *Collect adoption over a longer period of time.* As the current data set has been collected by periodically retrieving the songs listened to since the last update, this could be done by continuing this process and repeating the analyses at a later time.

3. *Explore different data sources.* In particular, other social networks could prove valuable resources. The wealth of information that is made public on these sites through profiles and connection lists could enable researchers to perform quantitative analyses on phenomena that until now could only be approached using surveys or secondary types of data. Marketers have certainly discovered this (Stanat, 2007), but a number of social scientists have also started examining the available data.

   In their soon-to-be-published study, Lewis et al. (2008) examined one college freshman class, as described by their Facebook profiles as well as additional information provided by the college. They not only collected cultural data on student tastes, but also gathered longitudinal data on the students' relationships within the confines of the college. In the latter, they use housing data as well as Facebook friends and pictures tagged by Facebook users to estimate the strength of each of the relationships.

Even with merely the data the university and the student provided, the analysis by Lewis et al. provided extensive information on:

- Social structure (network size, network density, heterogeneity, centrality)
- Population demographics, student diversity
- Subgroup differences across the three types of ties
- Cultural preferences and the intersection of tastes and ties

Their data was limited to a small user group, but it is not hard to imagine what increasing data sets and increased processing capacity can do for quantitative social research. This type of data could be very valuable in characterising users and creating a representative agent population. Not all social networks can be used to collect information about the adoption of new products (though some do, such as the adoption of Facebook applications). However, with this type of data, other social networks could still add significant value to the social sciences. It will be interesting to see how the observations on the effect of social networking sites presented in this thesis translate to other domains, such as the adoption of consumer durables.

4. *Develop new models for the adoption of music.* There has been work by both Delre (2007, pp. 87–109) and Broekhuizen & Delre (2008) on adapting the model used here to the motion picture industry. A similar approach could be taken to adapt the model to the dynamics of the music industry.

   In fact, commercial start-up 'Platinum Blue' has developed a properietary computer program that analyses the mathematical relationships among all of a song's structural components: melody, harmony, beat, tempo, rhythm, octave, pitch, chord progression, cadence, sonic brilliance, frequency, and so on. Based on these calculations, they claim to be able to predict the hit potential of a song with 80% accuracy.

   In fact, most hit songs fall in one of sixty 'hit clusters', grouping songs with similar characteristics. According to their findings, people who like the songs in that hit cluster are very likely to also enjoy new songs with similar characteristics, even if the sound is completely different (Gladwell, 2006).

   This would come to no surprise to those tracking Last.fm's competitors. One of the other major sites offering personalized music recommendations, Pandora, uses the results of their own 'Music Genome Project'. This project involves a group of musicians listening to every track in their library and identifying several hundred characteristics of each song (Pandora, 2008). These characteristics seem very similar to the components analysed by Platinum Blue's system.

Both companies claim that these characteristics can be used to predict whether someone will like a song. If this is indeed the case, these characteristics could be used as the basis of a new agent-based model, combining social processes with each user's preferences for each of these characteristics. Pandora, with its combination of large amounts of listener data and the Music Genome Project, is in an ideal position to build such a model, but other companies could move to collecting this data as well. This could be a very interesting development, indeed.

Apart from these model-related suggested, there are two topics that presented themselves over the course of the study:

1. *Improve validation methods.* Current validation methods (see subsection 2.1.3) are largely focused on calibration and model building. Both calibration and model building are vitally important, of course, but there could be a step missing. There seems to be no clear and standard method of determining whether a model is 'good enough'.

   Many other measures of fit have standard interpretations, such as Cohen's Kappa (Landis & Koch, 1977). These provide guidance in interpreting validation results. It seems the field of agent-based modelling would benefit greatly from a standardised method for determining the appropriate threshold value for R^2 or some other measure. Such a method would improve certainty about the validity of a model for a particular application.

2. *Re-examine the distinction between influentials and imitators.* Recent studies on social contagion (such as van den Bulte & Wuyts, 2007) have focused on a distinction between a small group of influentials and a larger majority of imitators. These studies suggest that there is a clear separation between the two groups: some people are influentials, some are imitators, and this rarely changes.

   The question is whether this clear separation will hold if the popularity of social networking sites continues to rise. The results presented in this thesis seem to suggest that heavy users of these sites are quick to respond to the music discoveries of their friends. However, it is doubtful that there are only a handful of users who discover of new music, who pass it on to the rest of the users. Since social networking sites focus on sharing, every user of these sites would seem more likely to influence friends with their music discoveries. Perhaps in future markets, users could more and more be considered to be at the same time influentials *and* imitators.

8.3 Other applications

A wide range of applications have been proposed for agent-based modelling research. For most of these, some type of micro-level data would improve the validity of the model and increase its usefulness. To give an example, among the many applications presented by Bonabeau (2002) are agent-based models of a supermarket and a department store. In these models, the behaviour of customers is simulated to improve store layout, reduce inventories and improve placement of salespeople. Since stores often collect very detailed data about consumer purchases, this is a prime example of a situation in which this data could be used to improve and extend existing agent-based models.

But consumer behaviour is not limited to store purchases. Recent simulation research at TNO (van Dam, 2008) has explored the consequences of decentralising energy generation, for example through the introduction of microCHP plants. These simulations were initialised using aggregate statistics of household composition and energy use. These simulations were very useful for exploring the consequences of different variants of decentralised energy generation. However, the simulation results were based on several assumptions that need to be validated before the conclusions can be used. Micro-level data set on household energy use would replace some of these assumptions with real-world data. As such, it would increase the usefulness of the conclusions as a basis for policy advice.

Bibliography

Alpaydin, E. (2004). *Introduction to Machine Learning*. Cambridge, MA, United States: The MIT Press.

Amaral, L. A. N., Scala, A., Barthélémy, M., & Stanley, H. E. (2000). Classes of small-world networks. *Proceedings of the National Academy of Sciences of the United States of America*, *97*(21), 11149–11152.

Arthur, W. B. (1991). Designing economic agents that act like human agents: A behavioral approach to bounded rationality. *The American Economic Review*, *81*(2), 353–359.

Arts, J. W. C., Frambach, R. T., & Bijmolt, T. H. A. (2006). What really drives innovation adoption by consumers? A meta-analysis on the antecedents of intention versus behaviour. Working paper.

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *289*, 509–512.

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

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, *99*(10, suppl. 3), 7280–7287.

Broekhuizen, T. L. J., & Delre, S. A. (2008). The effect of social influence on market inequalities in the motion picture industry. In *Proceedings of the European Marketing Academy Conference 2008*.

Carley, K. M. (1996). Validating computational models. *Working Paper Carnegie Mellon University*.

Chatterjee, R., & Eliashberg, J. (1990). The innovation diffusion process in the heterogeneous population: a micromodelling approach. *Management Science*, *36*, 1057–1079.

CNN (2006). If it's cool, creative and different, it's indie. Retrieved 12 July 2008. URL `http://www.cnn.com/2006/SHOWBIZ/Music/09/19/indie.overview/`

Delre, A. (2007). *The Effects of Social Networks on Innovation Diffusion and Market Dynamics*. Ph.D. thesis, University of Groningen, Groningen, The Netherlands.

Dosi, G., Fagiolo, G., & Roventini, A. (2006). An evolutionary model of endogenous business cycles. *Computational Economics*, *27*, 3–34.

Ellison, N., Steinfield, C., & Lampe, C. (2007). The benefits of Facebook "friends": Exploring the relationship between college students' use of online social networks and social capital. *Journal of Computer-Mediated Communication*, *12*(3).

Epstein, J. M. (1999). Agent-based computational models and generative social science. *Complexity*, *4*(5), 41–60.

Epstein, J. M., & Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge, MA, United States: The MIT Press.

Facebook (2008). Statistics. Retrieved 24 April 2008.
URL `http://www.facebook.com/press/info.php?statistics`

Fagiolo, G., Moneta, A., & Windrum, P. (2005). Empirical validation of agent-based models: a critical survey. *International Workshop on Agent-Based Modeling and Economic Policy Design.*. Germany: University of Bieleveld.

Garcia, R. (2005). Uses of agent-based modeling in innovation/new product development research. *Journal of Product Innovation Management*, *22*, 380–398.

Garcia, R., Rummel, P., & Hauser, J. (2007). Validating agent-based marketing models through conjoint analysis. *Journal of Business Research*, *60*, 848–857.

Gladwell, M. (2006). The formula. Retrieved 13 July 2008.
URL `http://www.gladwell.com/2006/2006_10_16_a_formula.html`

Golder, S., Wilkinson, D. M., & Huberman, B. A. (2006). Rhythms of social interaction: messaging within a massive online network.
URL `http://www.citebase.org/abstract?id=oai:arXiv.org:cs/0611137`

Guardiola, X., Díaz-Duilera, A., Pérez, C. J., Arenas, A., & Llas, M. (2002). Modeling diffusion of innovations in a social network. *Physical Review E*, *66*(026121), 1–4.

Hahn, M., Park, S., Krishnamurthi, L., & Zoltners, A. (1994). Analysis of new-product diffusion using a four-segment trial-repeat model. *Marketing Science*, *13*, 224–247.

Holland, J. H., & Miller, J. H. (1991). Artificial adaptive agents in economic theory. *The American Economic Review*, *81*(2), 365–370.

Jager, W. (2000). *Modelling Consumer Behaviour*. Ph.D. thesis, Centre for Environmental and Traffic Psychology, University of Groningen, Groningen, The Netherlands.

Jager, W. (2008). Personal communication.

Jain, D., Mahajan, V., & Muller, E. (1991). Innovation diffusion in the presence of supply restrictions. *Marketing Science*, *10*, 83–90.

Janssen, M. A., & Jager, W. (2003). Self organisation of market dynamics: Consumer psychology and social networks. *Artificial Life*, *9*(4), 343–356.

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, *33*, 159–174.

Last.fm (2008). Last.fm FAQ: How are the similar artists calculated? Retrieved 7 May 2008.
URL `http://www.last.fm/help/faq/?category=Artist+Pages#293`

Lewis, K., Kaufman, J., Gonzalez, M., Wimmer, A., & Christakis, N. (2008). Tastes, ties and time: a new social network dataset using Facebook.com. Unpublished paper.

Lucas, R. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, *1*, 19–46.

Mahajan, V., Muller, E., & Wind, J. (Eds.) (2000). *New Production Diffusion Models*. Boston, MA, United States: Kluwer Academic Publishers.

Malerba, F., Nelson, R., Orsenigo, L., & Winter, S. (1999). History-friendly models of industry evolution: the computer industry. *Industrial and Corporate Change*, *8*, 3–41.

Malerba, F., & Orsenigo, L. (2001). Innovation and market structure in the dynamics of the pharmaceutical industry and biotechnology: towards a history-friendly model. *Proceedings of the Conference in Honour of Richard Nelson and Sydney Winter, Aalborg, 12th-15th June 2001*.

Meade, N., & Islam, T. (2006). Modelling and forecasting the diffusion of innovation: A 25-year review. *International Journal of Forecasting*, *22*, 519–545.

Pandora (2008). The music genome project. Retrieved 21 July 2008.
URL `http://www.pandora.com/mgp.shtml`

Rogers, E. M. (1976). New product adoption and diffusion. *Journal of Consumer Research*, *2*, 290–301.

Ruiz, E. (2005). *Modelling Innovation Diffusion Patterns*. Ph.D. thesis, University of Groningen, Groningen, The Netherlands.

Ryan, B., & Gross, N. C. (1943). The diffusion of hybrid seed corn in two iowa communities. *Rural Sociology*, *8*(1), 15–24.

Schelling, T. C. (1978). *Micromotives and Macrobehaviour*. Toronto, Canada: George J. McLeod Ltd.

Stanat, M. (2007). Facebook & social networking in market research. Tech. rep., SIS International Research. Retrieved 29 February 2008.
URL `http://www.sisinternational.com/media/article.php/sis_facebook`

The Guardian (2008). Last.fm widgets boost user numbers.
URL `http://www.guardian.co.uk/media/2008/feb/28/web20.digitalmedia`

Twomey, P., & Cadman, R. (2002). Agent-based modelling of customer behaviour in the telecoms and media markets. *info*, *4*(1), 56–63.

Valente, T. W. (1996). Social network thresholds in the diffusion of innovations. *Social Networks*, *18*, 69–90.

van Dam, M. C. (2008). *Sustainable Power Play*. Master's thesis, University of Groningen.

van den Bulte, C. (2008). Personal communication.

van den Bulte, C., & Stremersch, S. (2004). Social contagion and income heterogeneity in new product diffusion: A meta-analytic test. *Marketing Science*, *23*(4), 530–544.

van den Bulte, C., & Wuyts, S. (2007). *Social Networks and Marketing*. Cambridge, MA, United States: Marketing Science Institute.

von Neumann, J., & Burks, A. W. (1966). *Theory of Self-Reproducing Automata*. Champaign, IL, United States: University of Illinois Press. Retrieved 7 July 2008.
URL `http://www.walenz.org/vonNeumann/`

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, *393*, 440–442.

Werker, C., & Brenner, T. (2004). Empirical calibration of simulation models. *Papers on Economics and Evolution*, *410*. Jena: Max Planck Institute for Research into Economic Systems.

Windrum, P., Fagiolo, G., & Moneta, A. (2007). Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, *10*(2).