UCLA

UCLA Previously Published Works

Title

Winner-take-all or long tail? A behavioral model of markets with increasing returns

Permalink

<https://escholarship.org/uc/item/0cs7q4q9>

Journal

System Dynamics Review, 32(3-4)

ISSN

0883-7066

Author

Lamberson, PJ

Publication Date

2016-07-01

DOI

10.1002/sdr.1563

Peer reviewed

MAIN ARTICLE

Winner-take-all or long tail? A behavioral model of markets with increasing returns

P.J. Lamberson^{a*}

Abstract

This paper develops a model of consumer choice that demonstrates why some markets with increasing returns converge to a winner-take-all outcome while many others have a power law market share distribution with a “long tail” of small-share products. The model takes the standard winner-take-all model of increasing returns and adds a simple behavioral assumption: when faced with complex choices, decision makers first quickly eliminate many of the available options using a simple heuristic before selecting from the remaining *feasible set*. We examine the market-level consequences of this model using an agent-based simulation. Under a wide range of parameters the model produces a power law share distribution. But when consumers have very large feasible sets the market converges to a winner-take-all outcome, and when consumers have very small feasible sets the model produces an evenly split market.

Copyright © 2017 System Dynamics Society

Syst. Dyn. Rev. **32**, 233–260 (2016)

Additional Supporting Information may be found online in the supporting information tab for this article.

Introduction

A variety of positive feedbacks ranging from network effects to economies of scale increase the likelihood that future consumers will purchase a product or adopt a technology as the size of the installed base of that product or technology grows (Sterman, 2000). Understanding the implications of these feedbacks for markets and firm strategy has been a fundamental contribution of the system dynamics methodology (e.g. Oliva *et al.*, 2003; Sterman *et al.*, 2007; Struben and Sterman, 2008).

In economics, these feedbacks are collectively known as *increasing returns*, and standard models predict that markets with increasing returns will converge to a winner-take-all (WTA) outcome (e.g. Arthur, 1989). Many empirical examples support this prediction, including the VHS versus Betamax and more recent Blu-Ray versus HD-DVD format wars (Arthur, 1990; den Uijl and de Vries, 2013), and early competition among nuclear reactor technologies (Cowan, 1990). In each of these cases a single technology eventually dominated the market (VHS, Blu-Ray, and light water reactors, respectively). But other markets with known positive feedbacks are split among multiple firms.

^a Communication Studies, University of California, Los Angeles

* Correspondence to: P. J. Lamberson. E-mail: lamberson@ucla.edu

Accepted by Markus Schwaninger, Received 26 April 2016; Revised 20 May and 16 November 2016; Accepted 19 January 2017

For example, the U.S. web browser market is shared by Google Chrome (34.7%), Microsoft Internet Explorer (28.3%), Safari (20.3%), Firefox (11%), and many others (Vaughn-Nichols, 2015). One site lists 40 browsers that are each estimated to account for at least one-tenth of 1 percent of the browser market (netmarketshare.com, 2015). Despite known cumulative advantages (Leskovec *et al.*, 2007), the markets for books, music, and movies all include many products that are each sold to a small number of people, but that in aggregate account for a large share of the market (Brynjolfsson *et al.*, 2003, 2006; Anderson, 2006). Rather than the WTA outcome that increasing returns theory predicts, these markets have a highly skewed or power law (PL) market share distribution, popularly known as a “long tail” (Anderson, 2006).ⁱ

This paper develops a model grounded in research on behavioral decision making that explains why some markets with positive feedbacks fit the WTA prediction, while others exhibit a PL. We argue that the standard model of increasing returns applies when there are few available choices and consumers are willing and able to consider all alternatives. The markets for VCRs, high-definition video disc players, and nuclear reactors (Arthur, 1990; Cowan, 1990) satisfy these conditions. We refer to this situation, and the corresponding model, as *simple increasing returns*. However, research in marketing and behavioral decision making demonstrates that when choosing among a large number of alternatives consumers employ a more complex decision-making algorithm; they first quickly eliminate many choices using a simple heuristic and then choose more carefully from the remaining *feasible set* (Simon, 1955; Einhorn, 1971; Payne, 1976; Hauser and Wernerfelt, 1990; Beach, 1993; Payne et al., 1993). In the model we propose, agents use such a two-stage decision rule, first limiting their choice to a feasible set and then selecting from among this set with a preference for more popular products. We call this model *complex increasing returns*.

We examine the market-level consequences of complex increasing returns through simulations and find that under an initial set of parameter choices the market share distribution converges to a PL. We then investigate the robustness of the PL result to an extensive set of parameter variations to better understand which features of the model drive the result. The model continues to produce a PL for a wide range of parameters, including different distributions of agent preferences and of available product attributes. However, the model produces different market share distributions with extreme distributions of agent feasible set sizes. By varying the distribution of feasible set sizes the model produces outcomes that range from the WTA outcome of simple increasing returns at the one extreme to the flat distribution predicted by standard models of differentiated product markets without positive feedbacks at the other (Lancaster, 1971).

Understanding the conditions under which increasing returns produce a WTA or PL market share distribution is more than an academic exercise. Strategy scholars recommend vastly different approaches in these two types of

markets. On the one hand, in WTA or winner-take-most markets, standard recommendations include discount pricing, heavy marketing, and rapid expansion to open a gap on competitors and “lock-in” the market (Arthur, 1996; Shapiro and Varian, 1999; Elberse, 2013). On the other hand, because each individual product in the long tail appeals to a limited audience, producers and retailers can ill afford to devote significant marketing resources to individual items. Rather, it is important to offer options that fit a wide range of consumer niches and to develop accurate “recommender systems” to help consumers find more obscure offerings that match their specific tastes (Anderson, 2006; Brynjolfsson *et al.*, 2006). The model in this paper suggests that individual feasible set sizes, a variable that has previously received little attention, is critical for understanding which of these strategy recommendations is appropriate in different positive feedback markets.

We further explore the implications of complex increasing returns by examining the impact of feasible set sizes on the predictability of market success. One key implication of standard models of simple increasing returns is that success is highly unpredictable; while one standard is guaranteed to eventually dominate the market, *which* will win is anybody’s guess. We find that with complex increasing returns the level of predictability depends strongly on the distribution of feasible set sizes. The results closely match experimental observations on consumer choice with social influence by Salganik *et al.* (2006).

We conclude by discussing how our model of complex increasing returns can help us to understand why, despite the increased availability of information on product popularity, the Internet is causing long tails to grow longer rather than markets to become more concentrated (Brynjolfsson *et al.*, 2003; Anderson, 2006).

Simple increasing returns and winner-take-all markets

We begin by describing the standard model of simple increasing returns and the WTA result. The underlying feedback structure of the model can be found in the textbook by Sterman (2000) and is reproduced in Figure 1.¹ Classic applications of the model include competition between VHS and Betamax video cassette formats, and the dominance of the QWERTY keyboard standard.

In the simplest version, a sequence of individuals $i = 1, 2, 3, \dots$ choose between two competing technologies, A and B . Each individual is either of type A or type B , and, all else equal, type A consumers prefer technology A to technology B , while type B consumers prefer technology B to technology A . In addition to the type related preferences, there is a positive feedback that

¹The model of simple increasing returns described here is based on the work of Arthur (1989). Arthur formulated his model as a discrete Markov process, but a continuous time version of the model can easily be implemented.

Fig. 1. The feedback structure of a standard network effects model (cf. Sterman, 2000, Figure 10–23)

increases the attractiveness of a technology as its market share increases. These preferences can be formalized using the utility function

$$U_i(x) = p_i(x) + \alpha n_x \quad (1)$$

Here, p_i is the type-based utility to individual i from choosing technology $x \in \{A, B\}$. For simplicity we assume that this utility is equal for all individuals of the same type, but later we show this assumption is unnecessary for the main implications. The fact that individuals prefer to purchase the technology that matches their type corresponds to the assumption that $p_i(A) > p_i(B)$ when i is of type A , and $p_i(B) > p_i(A)$ when i is of type B . The term αn_x captures the positive feedback, where n_x is the current installed base of technology x and α is simply a scaling constant that measures how much positive feedbacks matter relative to individual tastes in the consumer's decision. Typically, $\alpha \ll 1$. As enumerated by Sterman (2000, ch. 10), there are a variety of mechanisms that can lead individuals to prefer more popular products including network effects (Katz and Shapiro, 1985), conformity pressure (Jones, 1984; Bernheim, 1994), and social learning (Ellison, 1993).

The dynamics of this model are straightforward and can be most easily understood by focusing on the difference in accumulated installed base between the two technologies, $n_A - n_B$. Early in the process, when neither technology has accumulated a significant installed base, the p_i term dominates the utility

function, so individuals simply purchase according to their type. If the order in which consumers enter the market is random with respect to type, then the difference in accumulated installed base between the two technology types starts at zero and makes a random walk on the integers as illustrated in Figure 2. However, if one of the two technologies accumulates a sufficient advantage so that the positive feedback term αn_x outweighs the utility from type-based individual tastes $p_i(x)$, then both types of individuals will purchase the market leading technology. At this point we say the market is “locked-in” because all future consumers will purchase from the market leader. Specifically, letting $p_a(x)$ and $p_b(x)$ denote the type-based utility from technology x of type A and type B individuals, respectively, the market locks-in to technology A if

$$p_b(A) + \alpha n_A > p_b(B) + \alpha n_B \quad (2)$$

$$\alpha n_A - \alpha n_B > p_b(B) - p_b(A) \quad (3)$$

$$n_A - n_B > (1/\alpha)(p_b(B) - p_b(A)) \quad (4)$$

Similarly, the market locks-in to technology B if

$$n_B - n_A > (1/\alpha)(p_a(A) - p_a(B)) \quad (5)$$

The mathematical theory of random walks guarantees that with probability one $n_A - n_B$ will eventually cross one of these two thresholds and thus the market is guaranteed to lock-in to a “winner-take-all” outcome.

Fig. 2. The typical dynamics of Arthur’s model of simple increasing returns and lock-in from a single simulation run. The “A wins threshold” and “B wins threshold” are given by the inequalities in Eqs (4) and (5), respectively. [Colour figure can be viewed at wileyonlinelibrary.com]

Arthur (1989) goes further by proving that, with any number of technologies and a continuum of individual tastes p_i , one can replace the linear network effects term αn_x in Eq. (1) with a general increasing function $f(n_x)$ and, so long as there exists an $\epsilon > 0$ such that $f' > \epsilon$, the market will eventually converge to a WTA outcome with probability one. In other words, as long as each individual purchase has some non-vanishing positive effect on the attractiveness of a technology to future consumers, the market is guaranteed to converge eventually to a WTA outcome.

Complex increasing returns

The analysis in the previous section raises a question: if, under the assumptions of the standard model, increasing returns always lead to a WTA outcome, but we observe market sharing in settings where we know increasing returns operate, then what assumptions of the model are violated?

Here, we focus on the assumption that there are non-vanishing positive feedbacks for all products for all consumers.² Research in behavioral decision making, marketing, and consumer behavior finds that when faced with complex choices consumers rarely consider all possible alternatives. Instead, they employ a combination of heuristics, first coarsely categorizing products as acceptable or unacceptable, and then selecting from the acceptable options using a more sophisticated heuristic (Simon, 1955; Einhorn, 1971; Payne, 1976; Payne *et al.*, 1993). Because this research implies that for some consumers some products are unacceptable no matter how popular they become, positive feedbacks only apply to those products within a consumer's feasible set.

To formalize this behavioral observation, our model makes two changes to the standard model described in the previous section. First, we replace the individual preference term p_i with a standard spatial model of individual preferences. Each product x is described by a point, also denoted by x , in a characteristics space A , and each agent i has an ideal point p_i in A . For example, the characteristics space for laptop computers consists of all possible combinations of size, weight, processor speed, hard drive size, and memory, among other things. Agents are assumed to prefer products closer to their ideal point, where distance is measured by a function $d: A \times A \rightarrow [0, \infty)$. This is the core structure of the Lancaster spatial model (Lancaster, 1971), which also underlies many standard conjoint analysis models used to estimate consumer preferences in marketing (Green and Srinivasan, 1990). The resulting utility to agent i of purchasing product x is

²Lee *et al.* (2006) relax another assumption of Arthur's model and demonstrate the possibility of market sharing. Namely, consumers in their model only receive increasing returns from purchases made by their neighbors in a social network. However, Lee *et al.*'s model only considers two competing alternatives and thus cannot account for the long tail phenomena we seek to explain here.

$$U_i(x) = \alpha n_x(i) - d(p_i, x) \quad (6)$$

where as before α is a positive constant and $n_x(i)$ is the number of agents that purchased product x prior to agent i . Replacing Arthur's p_i with the distance term $-d(p_i, x)$ has no substantive impact on the model or the WTA result, and is merely a formal change that allows us to implement the second adaptation.

The second modification to the model is more significant. Namely, each agent in the complex increasing returns model has a *feasible set* of products $\mathcal{F}_i$, specified by the agent's ideal point p_i and a radius $r_i \in [0, \infty)$:

$$\mathcal{F}_i = \{x \mid d(p_i, x) \leq r_i\}$$

The first step in an agent's decision is to eliminate all products not in their feasible set. If the feasible set is non-empty, the agent selects the product in the feasible set that maximizes the combination of increasing returns and personal preferences in Eq. (6). If no products meet the agent's consideration requirements, he or she does not make a purchase (or, equivalently, chooses to purchase an outside good).

The model resonates with the shopping experience at a typical online retailer. For example, imagine a consumer shopping for a camera at an online electronics store. The website presents the consumer with a dizzying range of options. With a few clicks she quickly narrows the displayed options to those that meet her requirements. Once the choice set has been narrowed so that all (or most) of the remaining products are acceptable, she sorts them according to popularity to see which of the options that meet her needs are most popular with other customers. Figure 3 shows an example from a large electronics retailer's website. The consumer has limited the feasible set using a sequence of attribute requirements: Cameras & Camcorders > Digital Cameras > Point & Shoot Cameras > 15+ Megapixels > 5–6× > Image Stabilization > HD Movie Mode. The default option is to sort the remaining products according to popularity ("Best Selling").ⁱⁱ

The emergence of market structure

We examine the market structure that emerges from the individual choice model in the previous section using a computer simulation.³ Because of the large number of products in the markets we wish to simulate, the heterogeneity in agent ideal points and feasible sets, and the interaction between

³The simulation was implemented in the programming language R. Source code for the base case simulations can be found in the Online Supplementary Material.

Fig. 3. The website of a popular electronics retailer (the image has been altered to obscure the retailer's identity). The website allows the consumer to quickly filter her options until only those that meet her requirements remain. The default option is to sort the remaining products according to popularity ("Best Selling"). [Colour figure can be viewed at wileyonlinelibrary.com]

consumer purchase decisions, the model is implemented as an agent-based simulation.

An overview of the simulation structure is shown in Figure 4. Before running the simulation, the modeler must externally specify a number of model parameters listed in Table 1. We describe the particular choices for the parameters in Table 1 in subsequent sections and for now focus on the steps of the simulation assuming that the parameters are given.

Once the parameters are specified, a one-time initialization stage begins, during which a set of simulated products is created by drawing N_p points, $\{x_j\}_{j=1}^{N_p}$, from the product attribute distribution f_p . A vector $s = (0, \dots, 0)$ of length N_p is created to keep track of the sales of the N_p products.

Following this initialization stage, a loop runs over a sequence of simulated consumers $i = 1, 2, \dots, N_c$, keeping track of the sales s of the set of products $\{x_j\}_{j=1}^{N_p}$ as the simulated consumers make their selections based on the choice model described in the previous section. At each step in the loop a consumer i is created by drawing an ideal point $p_i \in A$ from the distribution f_c and a feasible set size r_i from the distribution g . If there are no products x_j with $d(x_j, p_i) \leq r_i$,

Fig. 4. Simulation overview

<i>Parametrization</i>	<ul style="list-style-type: none">Parameters from Table 1 set externally by modeler
<i>Initialization</i>	<ul style="list-style-type: none">Draw a collection of N_p product attribute points $x_j \in A$Create an initial sales vector $s = (0, \dots, 0)$ of length N_p
<i>Simulation loop</i>	
For consumers $i = 1, \dots, N_c$	<ul style="list-style-type: none">Draw a consumer ideal point p_iDraw a consumer feasible radius r_iIf $\mathcal{F}_i = \{x_j \mid d(x_j, p_i) \leq r_i\} = \emptyset$ move to step $i + 1$ in the loop i.e. no feasible products $\implies$ no purchaseElse find the product $x_j \in \mathcal{F}_i$ that maximizes $U_i(x_j) = \alpha s_j - d(p_i, x_j)$ and set $s_j = s_j + 1$
<i>Analysis</i>	Analyze the market shares $s / \sum_{j=1}^{N_p} s_j$

Table 1. Simulation parameters

Parameter	Description
N_p	Number of products
N_c	Number of consumers
A	Characteristics space
d	Distance metric
f_p	Product attribute distribution
f_c	Consumer ideal point distribution
g	Feasible set radii distribution
α	Utility weight on popularity

then the simulation moves on to the next consumer without recording a purchase. If the feasible set is non-empty then the simulated consumer i chooses the product x_j that maximizes the utility function

$$U_i(x_j) = \alpha s_j - d(p_i, x_j) \tag{7}$$

subject to the constraint that $d(p_i, x_j) \leq r_i$. The purchase is recorded by incrementing the sales vector s by one in the entry corresponding to the purchased product.

The simulation then moves on to the next simulated consumer and so on until the final consumer $i = N_c$ is reached. As the simulation proceeds, we keep track of the number of sales for each product. When the final consumer $i = N_c$ makes their choice the simulation terminates and we analyze the resulting distribution of sales. Although the simulation is dynamic and the market share distribution emerges over time (where “time” in the model is indexed by the sequence of consumer purchases), for this analysis we focus on the

distribution of market shares at $i=N_c$. Graphical examination of the market shares over time show that the market share distribution stabilizes at an effectively steady state well before the simulation terminates, so this analysis is equivalent to examining the steady state results of the dynamic process.

For each set of parameters chosen, we repeat the simulation many times and examine the distribution of outcomes. Below (“Analysis of the model”) we sweep across many values of the parameters to examine the sensitivity of these outcomes to the parameter choices. Thus the computational simulations allow us to build a theory of how the individual choice model described in the previous section relates to market-level outcomes, and how these market-level outcomes depend on specific parameters of the model (Miller and Page, 2007).

Benchmark simulations

Before examining the results of the full model, we consider two benchmark cases: no increasing returns and simple increasing returns. We then turn to the model of complex increasing returns and find that with a set of initial assumptions regarding agents and products, shown in Table 2, the model produces a PL. We vary the parameters to determine which assumptions drive the result. The distribution of market shares continues to follow a PL for many, but not all, parameter values.

The model is run 100 times with both no increasing returns and simple increasing returns (no feasible sets). All other parameters are as in Table 2 (the choice of base case parameters is discussed below, under “Complex increasing returns and the power law result”). These results will serve as useful benchmarks when we turn to the full model in subsequent sections. Figure 5 presents the resulting market share distribution from a representative run with no increasing returns in the left-hand panels and with simple increasing returns in the right-hand panels. The panels in the top row show the market shares for all products that earned at least one-tenth of 1 percent of the market, ordered by market share. The panels in the bottom row display the rank/frequency plots on log–log scales and the maximum likelihood estimate of a PL fit to the data.⁴

The first and third columns of Table 3 summarize the results. With no increasing returns agents simply choose the product closest to their ideal point. As expected, the 100 products split the market roughly equally: the mean market share is 0.01, the mean maximum share is 0.03 and on average 99.7 products earn a market share greater than 0.001.

These results echo those of classic models of competition among differentiated products by Hotelling (1929) and Lancaster (1971, 1979), which predict

⁴The rank/frequency plot is an unbiased estimator of the complementary cumulative distribution function, $F(x) = P(X \geq x)$, for the distribution from which the data are drawn. If the data are drawn from a PL, this plot approximates a straight line on log scales (for details see Appendix.)

Table 2. Base case parameter values

Parameter	Description	Base case value
N_p	Number of products	100
N_c	Number of consumers	10,000
A	Characteristics space	$[0, 10] \times [0, 10]$
d	Distance metric	Euclidean
f_p	Product attribute distribution	Uniform over A
f_c	Consumer ideal point distribution	Uniform over A
g	Feasible set radii distribution	Uniform $[0, \sqrt{50}]$ ¹
α	Utility weight on increasing returns	0.01

¹ $\sqrt{50}$ is the distance from the center of the attribute space to a far corner of the space. See section on “Complex increasing returns” for details.

Fig. 5. Market share distributions and rank/frequency plots from benchmark simulations. All other parameter values are shown in Table 2. Note that the vertical scale in the upper left panel differs from the vertical scale in the upper right panel

roughly uniform market share distributions. Unlike the WTA outcome under simple increasing returns, differentiation without popularity effects diverges from a PL in the opposite direction, producing a market that is “all tail” without the large share products at the head of the distribution characteristic of a PL. The market share distribution is determined entirely by the positions of agent ideal points and products in the characteristics space. Differences in

Table 3. Benchmark and base case simulation results

	No increasing returns	Complex increasing returns	Simple increasing returns
Products above min. share	99.7 (0.601)	48.6 (2.77)	10.7 (3.39)
Mean share	0.010 (0.000)	0.011 (0.000)	0.014 (0.001)
Median share	0.009 (0.000)	0.001 (0.000)	0.000 (0.000)
Max. share	0.030 (0.005)	0.380 (0.049)	0.962 (0.008)
KS statistic	0.189 (0.020)	0.056 (0.018)	0.318 (0.103)
Power law exponent (k)	1.60 (0.095)	1.68 (0.034)	1.92 (0.211)
Chose outside good	—	731 (37.5)	—

Means and standard deviations from 100 simulations. (Note: standard errors are not reported as these can be made arbitrarily small by increasing the number of simulations run.)

product attributes cause slight variation in market share—products that are clustered more closely together receive smaller shares of the market, while others that occupy less crowded areas of the attribute space have greater market shares—but these deviations from equal shares are small. Furthermore, with no increasing returns such a market share distribution would quickly dissolve if firms were allowed to strategically position their products. In that case, as argued by Hotelling (1929) and Lancaster (1971, 1979), firms would reposition themselves until all firms captured equal shares of the market. Thus consumer preferences alone are not a reasonable explanation for the broad empirical PL observation.

In the simple increasing returns case a single product monopolizes the market: the mean maximum share from the 100 simulations is 0.96, and on average only 10.7 products earn a market share greater than 0.001. These results demonstrate that adding product differentiation to a model of simple increasing returns is insufficient to explain the PL regularity.

The small share products shown in the upper right panel of Figure 9 are a relic of purchases made early on in the simulation run. If $D_{\max}$ is the maximum distance between any two points in the product characteristics space ($D_{\max} = \sqrt{10^2 + 10^2} = \sqrt{200}$ in these simulations), all agents will choose the same product once the best-selling product has accumulated $D_{\max}/\alpha$ more sales than any of its competitors. To see this, note that agent i will choose product x once

$$U_i(x) = an_x(i) - d(p_i, x) > an_y(i) - d(p_i, y) = U_i(y) \quad (8)$$

for all products y not equal to x . Some algebra shows that this is equivalent to

$$n_x(i) - n_y(i) > (d(p_i, x) - d(p_i, y))/\alpha \quad (9)$$

Since $(d(p_i, x) - d(p_i, y))$ is always less than or equal to $D_{\max}$, if $n_x(i) - n_y(i) > D_{\max}/\alpha$ then Eq. (9) is guaranteed to hold. The theory of random walks guarantees that one product will eventually accumulate such an advantage.

Thus only early adopters choose a product other than the eventual market leader. In each of the 100 runs of the simulation with simple increasing returns analyzed here, only the eventually dominant product captured any sales after the first 750 purchases. This also demonstrates that the WTA outcome with simple increasing returns is independent of parameter choices such as the distribution of agent ideal points or product characteristics, replicating the WTA outcome of the standard model of simple increasing returns.

Complex increasing returns and the power law result

We now turn to the results of the model under complex increasing returns. We begin by examining the simulation results under an initial “base case” set of parameters shown in Table 2, but *the focus of our analysis is not on any particular choice of parameters, but rather on how the output of the model varies with changes in the input parameters* as examined below (“Analysis of the model”). Our choice of base case parameters follows the framework of a “neutral model” developed and widely used in genetics and evolutionary biology, in which no particular individual has a significant a priori advantage (e.g. Hubbell, 2001). Later we examine the sensitivity of the model results to these parameter choices.

We now describe the base case parameters. The weight on product popularity $\alpha = 0.01$ only changes the speed with which the market reaches equilibrium. The number of consumers $N_c = 10,000$ is sufficiently large so that the market shares are stable well before the simulation is complete. The two-dimensional Euclidean topology with boundary for the attribute space A is chosen because it is computationally fast. Using a toroidal topology is examined below and does not significantly change the results. (As in the standard circular city model (Salop, 1979), a toroidal topology removes any potential boundary effects.) The uniform distribution of consumer ideal points f_c and product attributes f_p ensures that no particular product has a significant advantage a priori. The maximum feasible set radius $\sqrt{50}$ is the distance from the center of the attribute space to the farthest corner, and thus is the smallest radius for which it is possible that the entire attribute space could be considered by a single consumer. A consumer with the minimum feasible set radius of zero will only accept a product with attributes that exactly match those of her ideal product.

The simulation is run for the base case 100 times, and the center column of Table 3 presents the results. A stable (over time) and consistent (across simulation runs) market structure emerges. The mean market leader's share across the 100 simulations is 0.38—much larger than the mean maximum share of

0.03, with no increasing returns and significantly less than the mean maximum share of 0.96 observed under simple increasing returns. On average, 48.6 of the 100 products earned a share greater than or equal to 0.001, again placing the base case results between those obtained without increasing returns and with simple increasing returns (99.7 and 10.7, respectively).

Although the summary statistics for the complex increasing returns simulations tend to fall between those obtained in the two benchmark models examined in the previous section, the outcome is not simply an average of the two extremes. On the one hand, the model results differ from those predicted by standard models of simple increasing returns because a sizable share of the market is captured collectively by firms that individually only obtain small market shares. Unlike the small share firms in the tail of the distribution under simple increasing returns shown in Figure 9, these small but non-vanishing market shares are not relics of model “burn-in” that would ultimately vanish if we run the model long enough. As the number of purchases approaches infinity, these firms continue to hold on to their position in the tail of the distribution. On the other hand, the inequality of success between the outsize “hits” at the head of the market share distribution and the many small share products in the tail is far from the roughly uniform share distribution predicted by standard models of competition among differentiated products introduced by Hotelling (1929) and later extended by Lancaster (1971). Instead, the resulting distribution of market shares, as shown for a representative example in Figure 6, resembles a PL.

To test the hypothesis that the resulting market shares follow a PL, we fit a PL to the data using maximum likelihood estimation (MLE) and test the goodness of fit with the Kolmogorov–Smirnov (KS) test. The null hypothesis that the data are drawn from the PL distribution estimated using maximum likelihood cannot be rejected at the 0.1 significance level in 93 of the 100 simulations.⁵ Thus the model reproduces the widely observed PL pattern.

Analysis of the model

While some markets display PL distributions, not all do. For example, Figure 7 shows the market share distributions for deodorant and breakfast cereals, neither of which fit a PL (the KS test rejects PL at the 0.05 level in both cases). Ideally, the model will reveal what characteristics of a market and of consumer demand are more or less likely to produce this phenomenon. To investigate this we vary seven parameters of the model: four relating to the products and three relating to the agents. In the former category we vary the number of products, the number of characteristics, the topology of the

⁵Contrary to many common statistical tests, in the KS test the *inability* to reject the null hypothesis at a *higher* significance level corresponds to stronger evidence of fit (see Conover, 1980; Press *et al.*, 2002). The 0.1 significance level is the highest reported in the table published by Goldstein *et al.* (2004) used in this analysis, and thus provides the strongest available evidence for PL fit. For details see Appendix.

Fig. 6. Market share distribution (left panel) and rank/frequency plot (right panel) with line indicating maximum likelihood estimate of PL fit from a base case simulation. Parameter values are listed in Table 2 (KS statistic = 0.055). For details see Appendix

Fig. 7. Rank/frequency plots of the market share distribution for deodorant and breakfast cereals (Euromonitor International, 2007)

characteristics space, and the distribution of product characteristics. In the latter category we vary the distribution of agent ideal points, the information delay before agents perceive past sales, and the distribution of feasible set radii.

Table 4 reports whether or not the PL fits with these variations in parameter values (see Table A1 in Appendix for the goodness-of-fit statistics). In most cases a PL continues to fit the market share distribution. Figure 8 displays representative market share distributions for a subset of the parameter variations where a PL continues to fit the data. While in each case the market share distributions are slightly different, they all fit a PL well. Similar variation is also evident in the empirical data.

The number of products has little effect on the distribution of shares. The PL continues to fit when agents only consider a single characteristic and also when

Table 4. Power law fit under varying parameter assumptions

Parameter changed	Value	Power law fits?
Number of products	10	Yes
	50	Yes
	250	Yes
	500	Yes
Characteristics space dimensions	1	Yes
	3	Yes
	4	Yes
Characteristics space topology	Torus	Yes
Product characteristics distribution	Normal(5, 1)	Yes
	Normal (5, 3)	Yes
	Normal(5, 5)	Yes
Agent ideal point distribution	Normal(5, 1)	Yes
	Normal (5, 3)	Yes
	Normal (5, 5)	Yes
Information delay	10	Yes
	100	Yes
	250	Yes
	500	Yes
Feasible set radii distribution	Uniform [0, 13]	Yes
	Uniform [0, 11]	Yes
	Uniform [0, 9]	Yes
	Uniform [0, 5]	Yes
	Uniform [0, 3]	Yes
	Uniform [0, 1]	No
	Uniform [1, $\sqrt{50}$]	Yes
	Uniform [3, $\sqrt{50}$]	No
	Uniform [5, $\sqrt{50}$]	No

Here, a PL is said to fit the data if the KS test does not reject the PL hypothesis at the 0.1 significance level using the mean KS statistic from 20 runs of the simulation. Normal (μ, σ) denotes a truncated normal distribution with mean μ and standard deviation σ truncated to lie in the interval [0, 10]. Uniform [min, max] denotes a uniform distribution over the interval [min, max]. Distributions are for each dimension of the characteristics space. See Appendix for details.

they consider three or four. Changing the topology of the characteristics space to a torus (and thus removing the boundary) also does not change the market-level outcome. When either product characteristics or agent ideal points are distributed normally, the distribution continues to be well approximated by a PL, regardless of the variance in those distributions. Incorporating an information delay, so that agents do not immediately perceive changes in the market share distribution, has no effect on the long-run outcome. The robustness of the model

Fig. 8. Market share distributions and rank/frequency plots from representative simulation runs with four parameter variations. All other parameters are as in the base case

to variations such as changes in the distribution of product attributes and agent preferences increases its credibility. PL market share distributions have been observed across such a wide range of markets that any plausible explanation for the pattern must hold up under alternate assumptions. The results of the model are most sensitive to changes in the distribution of feasible set radii.

The distribution of feasible set sizes

In reality feasible set sizes vary across consumers within an industry, and the distribution of sizes varies across industries (Howard and Sheth, 1969; Belonax and Mittelsteadt, 1978; Hauser and Wernerfelt, 1990). In some industries, there may not be any consumers with a large feasible set. The phrase, “Give me the most popular car you’ve got” is likely a rarity in auto dealer showrooms. In other product markets, such as cell phones or portable music players, most consumers may simply want a product that works and is compatible with others but care little about the underlying technologies. In these markets most consumers have a large feasible set. The results of the model indicate that this variation in the specificity of consumer requirements can have dramatic consequences for market structure.

The distribution of feasible set sizes has a greater effect on the eventual market form than either the distribution of product characteristics or the distribution of agent preferences. When the minimum feasible set radius is large (three

or five), the top few products are unusually successful; that is, they win a higher share of the market than predicted by the maximum likelihood estimated PL distribution.

The left-hand panel of Figure 9 shows a representative rank/frequency plot with agent radii distributed $U[5, \sqrt{50}]$. These results more closely resemble those under simple increasing returns than the base case of complex increasing returns. The mean shares for the top three products respectively are 0.7, 0.16 and 0.1. In this case, the average median number of products in an agent's feasible set is 62, so for a majority of the agents most of the products are acceptable, making the consumers' decisions more similar to simple increasing returns where all products are acceptable for all agents. This could correspond to an industry in which compatibility and other network externalities are more important to consumers than specific product characteristics, such as the market for portable digital music players, in which the top three firms have market shares 74 percent, 9 percent, and 3 percent (Elmer-DeWitt, 2007). Clearly, if the minimum feasible set radius was sufficiently large so that all feasible sets contained all products, complex increasing returns would replicate the model of simple increasing returns and converge to a WTA outcome. Thus simple increasing returns are an extreme case of complex increasing returns in which all agents have very large feasible sets.

As shown in the right-hand panel of Figure 9, when the maximum feasible set radius is small, the distribution of shares looks similar to the results without increasing returns (shown in the right-hand panel of Figure 5). In this case, where the maximum feasible set radius is one, the average agent's feasible set contains only one product, and thus most agents choose based on personal preferences alone without being influenced by product popularity. Deodorant and breakfast cereals have similar market share distributions, shown in

Fig. 9. Rank/frequency plots from a representative simulation with parameter values as in Table 2 with feasible set radii distributed $U[5, \sqrt{50}]$ in the left-hand panel and with feasible set radii distributed $U[0, 1]$ in the right-hand panel

Figure 7. The implication is that for these products most consumers consider very few of the available options or increasing returns play no role.

While this is not an exhaustive sweep of all potential distributions of feasible set sizes, some general features of the model are evident. When agents have specific tastes and pay little attention to the choices of others, the market is split roughly equally among all products. When the majority of agents find most of the products acceptable, the results are closer to the WTA distribution observed under simple increasing returns. For intermediate distributions of agent feasible sets the market share distribution follows a PL.

As these results demonstrate, despite the assumption that consumers value previous sales equally across all of the simulations (i.e. α is constant), a variety of market structures arise depending on the restrictiveness of consumer requirements. Markets that consist only of a long tail with no massive hits or only of hits with no long tail lie at the extreme ends of a spectrum controlled by the size of consumer feasible sets. For a wide range of intermediate distributions of agent feasible sets the market share distribution follows a PL, and this outcome is robust to other variations in the model.

Predicting market success

Anecdotally, markets with increasing returns are notoriously unpredictable (Gladwell, 2006; Watts, 2007), and competing in these markets has been likened to high-stakes gambling (Arthur, 1996; Watts and Hasker, 2006). However, empirically measuring unpredictability is difficult since history only runs once. The simulation approach allows us to examine the predictability of markets with complex increasing returns by running the simulation multiple times and examining the consistency of the outcomes. We run the model 100 times with a single set of products (with the base case assumptions on the distribution of product characteristics shown in Table 2) under each of three conditions: no increasing returns, complex increasing returns, and simple increasing returns; and then compute the measure of unpredictability used by Salganik *et al.* (2006).⁶ The unpredictability for each product is given by

$$u_i = \frac{\sum_{j=1}^{100} \sum_{k=j+1}^{100} |m_{i,j} - m_{i,k}|}{\binom{100}{2}} \quad (10)$$

where $m_{i,j}$ is the market share of product i in repetition j . The total number of pairs of repetitions is $\binom{100}{2}$, so Eq. (10) simply gives the average difference in

⁶Salganik *et al.* (2006) conduct an Internet-based experiment to examine the effect of social influence on consumer behavior. They compare the downloads of a collection of songs from a control group where individuals are given no information about other consumers' choices, and a treatment group, where individuals are shown each product's number of previous downloads.

Fig. 10. Unpredictability with no increasing returns, complex increasing returns, and simple increasing returns from 100 simulation runs. Differences in the unpredictability between each pair of conditions are significant ($p < 0.0001$)

market share for product i over all pairs of repetitions. The unpredictability for each returns condition is the average over all products. Because the only stochastic variables in this exercise are the consumer ideal points and feasible sets, this quantity captures how unpredictable the market outcomes are even when the overall distribution of consumer preferences and product attributes remains constant.

Figure 10 plots the unpredictability for each of the three conditions. Without increasing returns, any unpredictability results from randomness in agent ideal points. After 10,000 agents the effect of this randomness on market shares is minimal. Complex and simple increasing returns have higher levels of unpredictability. These results mirror those obtained experimentally by Salganik *et al.* (2006), who found unpredictability to be significantly higher with social influence than without and higher when social influence was more salient.⁷

As discussed above, by varying the distribution of agent feasible set radii, complex increasing returns produces results ranging between those expected from choice without increasing returns to those of simple increasing returns. Similarly, the corresponding level of unpredictability moves from low values when feasible sets are small to higher values as feasible sets become large. To demonstrate this, we run the same experiment under complex increasing returns with three additional distributions of feasible set sizes: $U[0,1]$, $U[0,3]$ and $U[0,13]$. Figure 11 plots the unpredictability with these three distributions of feasible sets along with the base case (feasible sets distributed $U[0, \sqrt{50}]$) and the no increasing returns and simple increasing returns cases. Clearly, unpredictability increases as feasible sets become larger.

⁷Salganik *et al.* (2006) have two social influence conditions: one in which they show subjects the number of previous purchases by other participants, and a second in which subjects are shown the number of previous purchases and the options are listed in order by the number of previous purchases.

Fig. 11. Unpredictability with complex increasing returns and different feasible set distributions along with no increasing returns and simple increasing returns from 100 simulation runs. $U[0, \sqrt{50}]$ is the base case. All differences between levels of unpredictability are significant ($p < 0.0001$)

For marketers, these results should raise a red flag. In any market where social influence plays a role, the specificity of consumer requirements should be a part of the marketing analysis. When consumers have larger feasible sets, effort spent on measuring and meeting consumer preferences may well be wasted if by chance a competitor gains an early advantage that overwhelms the benefits of producing a well-positioned product. When consideration sets are small, traditional methods of determining consumer preferences, such as conjoint analysis, are more likely to be successful, but as consideration sets grow larger it becomes more important that forecasting efforts consider the impact of social influence on market success.

Discussion

By adding an important behavioral assumption—the use of feasible sets—this model extends the theory of increasing returns to markets where consumers choose among many products and no single product comes to dominate the market. Research from marketing, consumer behavior, and psychology demonstrates that consumers often use feasible sets when making complex choices (Einhorn, 1971; Payne, 1976; Olshavsky, 1979; Payne *et al.*, 1993). This paper extends this research on individual use of the feasible set heuristic to examine the market-level implications. The findings demonstrate that market structure depends on the size of consumers' feasible sets. When consumers consider a wide variety of products, the market reflects the WTA outcome that we tend to associate with increasing returns. When consumers limit their choice, and thus hold smaller feasible sets, the market reflects the flat distribution of the traditional spatial model. Between these two extremes the model most often produces a PL.

Although a variety of statistical models generate PLs (Mitzenmacher, 2004), this model gives us insight into a microlevel choice process that gives rise to the macrolevel pattern. That insight sheds light on a puzzle of the Internet economy: information on product popularity is more accessible to consumers

than ever before and yet, contrary to theoretical predictions, markets are becoming less concentrated (Brynjolfsson *et al.*, 2003; Anderson, 2006). While the Internet provides more information about product popularity, it also gives consumers more choices and more information about product attributes (Brynjolfsson *et al.*, 2003, 2006; Anderson, 2006). By making product characteristics accessible, the Internet facilitates the screening process, enabling “the formation of consideration sets that include only those few alternatives best suited to a consumer’s personal taste” (Alba *et al.*, 1997). In an experiment, Belonax and Mittelsteadt (1978) find that individuals consider fewer products when information about more product characteristics is available. The model of complex increasing returns in this paper predicts that the resulting narrowing of individuals’ feasible sets will lead to less concentrated market share distributions with more niche products.

One next step suggested by this paper is to empirically verify the predicted relationship between feasible set sizes and market share distributions. Specifically, when all feasible sets contain most of the available products, markets will be dominated by a single product; when all feasible sets contain very few products, markets will be evenly split; and when feasible sets are more variable, the market share distribution resembles a PL. At present, little data on feasible set sizes is available. The best available data come from automobiles and consumer packaged goods where feasible set sizes are small and, as the model predicts, market share distributions are relatively flat (Hauser and Wernerfelt, 1990). But the results of this model suggest that feasible set sizes should be an important variable to examine in future empirical research on consumer choice. A laboratory experiment, similar to the design employed by Salganik *et al.* (2006), but with an added feasible set component, could also be used to test the model predictions.

A second area for future research involves expanding the model to include more detail, such as social network structure on the demand side as considered by Lee *et al.* (2006), and product positioning dynamics on the supply side. Most studies of strategies for competing in markets with increasing returns have focused on pricing (Arthur and Rusczyński, 1992; Fudenberg and Tirole, 2000) and compatibility choices (Katz and Shapiro, 1985, 1986). By introducing horizontal differentiation among products, the complex increasing returns model raises a new dimension of firm strategy to be examined: the choice of product attributes. Some authors argue that firm strategies exhibit significant “herding behavior” or “bandwagon effects” (Abrahamson and Rosenkopf, 1993; Haveman, 1993; Kennedy, 2002; Scharfstein and Stein, 1990), but an imitative strategy in a complex increasing returns market is unlikely to succeed because consumers will always choose the more popular of two similar products. Instead, a firm is better off distinguishing itself in order to capture consumers that do not include more popular competitors in their feasible set. Even in this simplified setting, an optimal product positioning strategy is difficult to determine.

Notes

- i. For empirical evidence of the PL market share regularity see Sutton (1997), Adamic and Huberman (2000), Huberman (2001), Levene *et. al.* (2001), Brynjolfsson *et. al.* (2003), Chevalier and Goolsbee (2003), and Kohli and Sah (2006). While some of these authors identify what we view as key ingredients in producing the PL regularity, they stop short of developing a full theory behind the observation. For example, Brynjolfsson *et. al.* (2006) point to variety in consumer preferences along with new Internet channel search tools, such as recommender systems and online product reviews, as key drivers of the long tail phenomenon. Similarly, Brynjolfsson *et. al.* (2011) find that reduced search costs are associated with a reduction in market share concentration and a shift towards a long tail. The advantage of our approach, and model building in general, is that it allows us to determine if indeed these factors alone can account for the observed PL pattern and, further, to test the sensitivity of that prediction to our modeling assumptions.
- ii. Note that while determining which products do and do not belong in a consumer's feasible set requires knowing the attributes of each product, in many settings this work is accomplished by an automated system. Even in the absence of an automated screening system, a consumer can often quickly create a feasible set without examining each product individually. For example, an individual shopping for a book on C++ programming in a bricks and mortar bookstore can head to the computer programming section of the store and safely assume that the books outside of that section are outside of her feasible set.

Acknowledgement

The author wishes to acknowledge the Ford Motor Company for financial support of this research through the University of Michigan-Ford alliance. Also, the author thanks Shayne Gary, Georgia Kernell, Sam Kernell, Scott Page, Irv Salmeen, and John Sterman for comments on earlier drafts of the paper. Brian Arthur, Robert Axtell, Mark Everson, Bryan Goodman, Mark Newman, Gint Puskorius, Carl Simon, John Sullivan, and many participants at the 2006 Santa Fe Institute Business Network Meeting provided valuable insights that helped in the shaping of this research.

References

- Abrahamson E, Rosenkopf L. 1993. Institutional and competitive bandwagons: using mathematical modeling as a tool to explore innovation diffusion. *Academy of Management Review* **18**: 487–487.
- Adamic LA, Huberman BA. 2000. The nature of markets in the world wide web. *Quarterly Journal of Electronic Commerce* **1**(1): 5–12.

- Alba J, Lynch J, Weitz B, Janiszewski C, Lutz R, Sawyer A, Wood S. 1997. Interactive home shopping: consumer, retailer, and manufacturer incentives to participate in electronic marketplaces. *Journal of Marketing* **61**: 38–53.
- Anderson C. 2006. *The Long Tail*. Hyperion: New York.
- Arthur WB. 1989. Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal* **99**: 116–131.
- Arthur WB. 1990. Positive feedbacks in the economy. *Scientific American* **262**: 92–99.
- Arthur WB. 1996. Increasing returns and the new world of business. *Harvard Business Review* **74**: 100–111.
- Arthur WB, Rusczyński A. 1992. Dynamic equilibria in markets with a conformity effect. *Archives of Control Sciences* **37**: 7–31.
- Beach LR. 1993. Broadening the definition of decision making: the role of prechoice screening of options. *Psychological Science* **4**(4): 215–220.
- Belonax J, Mittelsteadt R. 1978. Evoked set size as a function of number of choice criteria and information variability. *Advances in Consumer Research* **5**(1): 48–51.
- Bernheim BD. 1994. A theory of conformity. *Journal of Political Economy* **102**(5): 841–877.
- Brynjolfsson E, Hu Y, Smith MD. 2003. Consumer surplus in the digital economy: estimating the value of increased product variety at online booksellers. *Management Science* **49**(11): 1580–1596.
- Brynjolfsson E, Hu Y, Smith MD. 2006. From niches to riches: the anatomy of the long tail. *Sloan Management Review* **47**(4): 67–71.
- Brynjolfsson E, Hu Y, Simester D. 2011. Goodbye pareto principle, hello long tail: the effect of search costs on the concentration of product sales. *Management Science* **57**(8): 1373–1386.
- Chevalier J, Goolsbee A. 2003. Measuring prices and price competition online: Amazon.com and barnesandnoble.com. *Quantitative Marketing and Economics* **1**: 203–222.
- Clauset A, Shalizi C, Newman M. 2009. Power-law distributions in empirical data. *SIAM Review* **51**(4): 661–703.
- Conover WJ. 1980. *Practical Nonparametric Statistics*. Wiley: New York.
- Cowan R. 1990. Nuclear power reactors: a study in technological lock-in. *Journal of Economic History* **50**(3): 541–567.
- den Uijl S, de Vries HJ. 2013. Pushing technological progress by strategic manoeuvring: the triumph of Blu-ray over HD-DVD. *Business History* **55**(8): 1361–1384.
- Einhorn HJ. 1971. Use of nonlinear, noncompensatory models as a function of task and amount of information. *Organizational Behavior and Human Performance* **6**: 1–27.
- Elberse A. 2013. *Blockbusters: Hit-Making, Risk-Taking, and the Big Business of Entertainment*. Macmillan: New York.
- Ellison G. 1993. Learning, local interaction, and coordination. *Econometrica* **61**: 1047–1047.
- Elmer-DeWitt P. 2007. Apple 2.0. Available: http://blogs.business2.com/apple/2007/04/sansa_connect_a.html[15 February 2017].
- Euromonitor International. 2007. Global market information database.
- Fudenberg D, Tirole J. 2000. Pricing a network good to deter entry. *Journal of Industrial Economics* **48**(4): 373–390.
- Gladwell M. 2006. The formula. *New Yorker* **16**: 138–149.

- Goldstein ML, Morris SA, Yen GG. 2004. Problems with fitting to the power-law distribution. *European Physical Journal B* **41**: 255–258.
- Green P, Srinivasan V. 1990. Conjoint analysis in marketing: new developments with implications for research and practice. *Journal of Marketing* **54**(4): 3–19.
- Hauser JR, Wernerfelt B. 1990. An evaluation cost model of consideration sets. *Journal of Consumer Research* **16**: 393–408.
- Haveman H. 1993. Follow the leader: mimetic isomorphism and entry into new markets. *Administrative Science Quarterly* **38**(4): 593–627.
- Hotelling H. 1929. Stability in competition. *Economic Journal* **39**(153): 41–57.
- Howard JA, Sheth JN. 1969. *The Theory of Buyer Behavior*. Wiley: New York.
- Hubbell SP. 2001. *The Unified Neutral Theory of Biodiversity and Biogeography*. Princeton University Press: Princeton, NJ.
- Huberman BA. 2001. *The Laws of the Web: Patterns in the Ecology of Information*. MIT Press: Cambridge, MA.
- Jones SR. 1984. *The Economics of Conformism*. Blackwell: Oxford.
- Katz ML, Shapiro C. 1985. Network externalities, competition, and compatibility. *American Economic Review* **75**(3): 424–440.
- Katz ML, Shapiro C. 1986. Technology adoption in the presence of network externalities. *Journal of Political Economy* **94**(4): 822–841.
- Kennedy R. 2002. Strategy fads and competitive convergence: an empirical test for herd behavior in prime-time television programming. *Journal of Industrial Economics* **50**(1): 57–84.
- Kohli R, Sah R. 2006. Some empirical regularities in market shares. *Management Science* **52**(11): 1792–1798.
- Lancaster KJ. 1971. *Consumer Demand: A New Approach*. Columbia University Press: New York.
- Lancaster K. 1979. *Variety, Equity, and Efficiency*. Columbia University Press: New York.
- Lee E, Lee J, Lee J. 2006. Reconsideration of the winner-take-all hypothesis: complex networks and local bias. *Management Science* **52**(12): 1838–1848.
- Leskovec J, Adamic LA, Huberman BA. 2007. The dynamics of viral marketing. *ACM Transactions on the Web* **1**(1): Article No. 5
- Levene M, Borges J, Loizou G. 2001. Zipf's law for web surfers. *Knowledge and Information Systems* **3**(1): 120–129.
- Miller JH, Page SE. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press: Princeton, NJ.
- Mitzenmacher M. 2004. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics* **1**(2): 226–251.
- netmarketshare.com. 2015. Desktop browser version market share. Available: <https://www.netmarketshare.com>. accessed date 25 March 2017.
- Newman MEJ. 2005. Power laws, Pareto distributions and Zipf's law. *Contemporary Physics* **46**(5): 323–351.
- Oliva R, Sterman JD, Giese M. 2003. Limits to growth in the new economy: exploring the “get big fast” strategy in e-commerce. *System Dynamics Review* **19**(2): 83–117.
- Olshavsky RW. 1979. Task complexity and contingent processing in decision making: a replication and extension. *Organizational Behavior and Human Performance* **24**(3): 300–316.

- Payne JW. 1976. Task complexity and contingent processing in decision making: an information search and protocol analysis. *Organizational Behavior and Human Performance* **16**: 366–387.
- Payne JW, Bettman JR, Johnson EJ. 1993. *The Adaptive Decision Maker*. Cambridge University Press: Cambridge, U.K.
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP. 2002. *Numerical Recipes in C*. Cambridge University Press: Cambridge, U.K.
- Salganik M, Dodds P, Watts D. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* **311**: 854–856.
- Salop SC. 1979. Monopolistic competition with outside goods. *Bell Journal of Economics* **10**(1): 141–156.
- Scharfstein D, Stein J. 1990. Herd behavior and investment. *American Economic Review* **80**(3): 465–479.
- Shapiro C, Varian HR. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press: Cambridge, MA.
- Simon HA. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* **69**(1): 99–118.
- Sterman J. 2000. *Business Dynamics*. Irwin/McGraw-Hill: New York.
- Sterman JD, Henderson R, Beinhocker ED, Newman LI. 2007. Getting big too fast: strategic dynamics with increasing returns and bounded rationality. *Management Science* **53**(4): 683–696.
- Struben J, Sterman JD. 2008. Transition challenges for alternative fuel vehicle and transportation systems. *Environment and Planning B: Planning and Design* **35**(6): 1070–1097.
- Sutton J. 1997. Gibrat's legacy. *Journal of Economic Literature* **35**: 40–59.
- Vaughn-Nichols SJ. 2015. Most popular U.S. web browsers, according to the federal government. Available: <http://www.zdnet.com/article/the-most-u-s-popular-web-browsers/> [15 February 2017].
- Watts DJ. 2007. The accidental influentials. *Harvard Business Review* **85**(2): 22–23.
- Watts DJ, Hasker S. 2006. Marketing in an unpredictable world. *Harvard Business Review* **84**(9): 25–30.

Appendix: Goodness-of-fit statistics

A random variable X has a PL distribution if its probability density function is of the form

$$f(x) = Cx^{-k}$$

where C and k are positive constants. Such a random variable is much less likely to take on large values than small values. When plotted on log–log scales, the probability density function is a straight line with slope $-k$ and intercept $\log C$. Because the density function $f(x)$ diverges as x approaches 0, any variable that follows a PL distribution can only do so above some minimum value strictly greater than zero. The researcher must make a judgment for the minimum value over which the PL holds (Newman, 2005). Also, for this

reason, the cumulative distribution function is measured from the right as opposed to the left, as is more standard for other distributions. This is sometimes referred to as the complementary cumulative distribution function.

The cumulative distribution function $F(x)$ is given by

$$F(x) = \text{probability}(X \geq x) = \int_x^{\infty} f(y) dy = \int_x^{\infty} Cy^{-k} dy = \frac{C}{k-1} x^{-(k-1)}$$

This is the same functional form as the probability density function, and thus also follows a straight line when plotted on log scales, but with a slope of $-(k-1)$ and intercept $\log \frac{C}{k-1}$. Many researchers make use of this fact to fit a PL to data (Goldstein *et al.*, 2004). An unbiased estimator of the cumulative distribution function is given by $S(x)$ = the fraction of data points greater than or equal to x (often called the empirical distribution function (Conover, 1980)). Given data $x_1 \dots x_n$, a graph of $S(x_i)$, commonly called a rank/frequency plot (Newman, 2005), can be made by simply plotting the rank of x_i divided by n (where the highest value of the x_i is ranked one) against x_i . If the x_i are drawn from a PL distribution, the rank/frequency plot will follow a straight line on log scales with the same slope and intercept as the cumulative distribution function (for an example, see Figure 6). Using a linear regression of $S(x_i)$ on x_i , one can estimate the parameters k and C . Although this method is widely used, it is known to be biased, and consequently the R^2 from this regression is not a recommended measure for goodness of fit (Goldstein *et al.*, 2004; Clauset *et al.*, 2009).

An unbiased estimator of the PL parameters can be obtained using MLE:

$$k = 1 + n \left(\sum_{i=1}^n \ln \frac{x_i}{x_{\min}} \right)^{-1},$$

$$C = (k-1)x_{\min}^{k-1}$$

where $x_i, \dots, x_n$ are the data, and $x_{\min}$ is the minimum of the x_i (Newman, 2005). The Kolmogorov–Smirnov (KS) test can be used to measure the goodness of fit for the estimated PL distribution.

The KS test is used to test the goodness of fit of data $x_1, \dots, x_n$ to a given distribution f . The null hypothesis is that the data are drawn from the given distribution. The test statistic D is defined by $D = \sup_x |F(x) - S(x)|$, where $F(x)$ is the cumulative distribution function for f and $S(x)$ is the empirical distribution function of the data. Comparing this statistic to a table of critical values, one can determine the significance level at which to reject the null hypothesis that the data were drawn from f . Note that for the KS test statistical significance indicates a lack of fit between the data and a PL. Thus the *inability* to reject the null hypothesis at a *higher* significance level corresponds to stronger evidence of fit.

The table for the test differs when parameters of the distribution f are estimated from the data $x_1, \dots, x_n$. We use the KS test table reported in

Goldstein *et al.* (2004) for MLE estimates of PL parameters. For further details see Goldstein *et al.* (2004), Press *et al.* (2002), or Conover (1980). In Table A1 we report both the KS statistic of the MLE and the R^2 from linear regression of the empirical distribution function to allow for comparison with empirical papers in which the latter method is used.

Table A1. Power law fit under varying parameter assumptions

Parameter changed	Value	KS statistic	$P \geq \min$	R^2
Number of products	10	0.146 (0.042)	10 (0.6)	0.912 (0.044)
	50	0.060 (0.016)	37 (2.2)	0.984 (0.006)
	250	0.066 (0.018)	59 (2.9)	0.992 (0.003)
	500	0.078 (0.026)	60 (4.2)	0.990 (0.003)
Characteristics space dimensions	1	0.104 (0.027)	21 (1.1)	0.988 (0.003)
	3	0.043 (0.014)	70 (3.8)	0.990 (0.003)
	4	0.047 (0.010)	82 (3.8)	0.984 (0.004)
Characteristics space topology	Torus	0.061 (0.019)	48 (2.4)	0.993 (0.003)
Product characteristics distribution	Normal (5, 1)	0.114 (0.043)	17 (1.3)	0.984 (0.007)
	Normal (5, 3)	0.065 (0.025)	44 (3.2)	0.990 (0.003)
	Normal (5, 5)	0.050 (0.010)	44 (3.5)	0.991 (0.003)
Agent ideal point distribution	Normal (5, 1)	0.112 (0.035)	15 (2.0)	0.961 (0.014)
	Normal (5, 3)	0.056 (0.01)	45 (2.6)	0.988 (0.004)
	Normal (5, 5)	0.062 (0.021)	49 (2.7)	0.991 (0.003)
Information delay	10	0.059 (0.022)	49 (2.5)	0.992 (0.003)
	100	0.051 (0.012)	53 (2.7)	0.992 (0.003)
	250	0.048 (0.01)	58 (3.3)	0.992 (0.003)
	500	0.042 (0.008)	65 (3.1)	0.993 (0.002)
Feasible set radii distribution	Uniform [0, 13]	0.087 (0.032)	38 (2.0)	0.983 (0.005)
	Uniform [0, 11]	0.065 (0.025)	41 (3.1)	0.987 (0.005)
	Uniform [0, 9]	0.055 (0.012)	44 (2.8)	0.991 (0.003)
	Uniform [0, 5]	0.046 (0.012)	57 (2.7)	0.985 (0.004)
	Uniform [0, 3]	0.061 (0.006)	69 (2.7)	0.945 (0.008)
	Uniform [0, 1]	0.207 (0.007)**	91 (2.3)	0.643 (0.020)
	Uniform [1, $\sqrt{50}$]	0.055 (0.016)	32 (1.6)	0.986 (0.005)
	Uniform [3, $\sqrt{50}$]	0.207 (0.058) [#]	15 (2.6)	0.946 (0.022)
	Uniform [5, $\sqrt{50}$]	0.290 (0.043)*	12 (2.9)	0.922 (0.023)

Means and standard deviations from 20 repetitions of the simulation with the given parameter values. All other parameters are as in Table 2. The column $P \geq \min$ is the number of products obtaining a market share greater than or equal to 0.001. Normal distributions for agent ideal points and product characteristics were truncated to lie in the interval [0, 10]. Distributions are for each dimension of the characteristics space. The following symbols indicate the level of significance p at which the null hypothesis (the data are drawn from the MLE PL distribution) can be rejected using the mean value of the KS statistic:

[#] $p < 0.1$;

* $p < 0.01$;

** $p < 0.001$.